Estimating the Utilisation of Key Licence-Exempt Spectrum Bands

Final Report

Issue 3
April 2009

Prepared By:

A J Wagstaff
System Design Authority

Approved By:

N P Merricks
Project Manager

Authorised By:

M J Ashman
Division Manager

Mass Consultants Limited
Grove House, Rampley Lane
Little Paxton, St Neots, Cambs PE19 6EL
National Tel: 01480 222600 Fax: 01480 407366
International Tel: +44 1480 222600 Fax: +44 1480 407366
E-mail: systems@mass.co.uk

ABSTRACT

A survey of IEEE 802.11b/g 'WiFi' usage has been carried out at various urban locations in the UK. This has revealed a wide variety of problems encountered by users, many of which are due to causes other than spectrum issues. These include problems with the wired Internet and device configuration errors.

Where the users' problems are spectrum-related they tend to be due to interference between devices in the 2.4 GHz ISM band rather than congestion, as was initially believed. Such problems are common but geographically dispersed. In the centre of London, however, it has been shown that the demands on this band are much higher than other locations surveyed and users are experiencing the combined effects of interference and congestion.

The problems of interference between different types of radio device in the 2.4 GHz band lead us to conclude that a certification scheme is highly desirable. Some non-WiFi devices already sport 'WiFi-friendly' claims on their datasheets. We propose extending this concept to a '2.4 GHz friendly' logo, which would help drive acceptance of innovative technologies in this band.

ACKNOWLEDGEMENTS

Our thanks are extended to all the many people and organisations who contributed to this project. In particular we would like to thank Intellect, OpenSpectrum, Cochrane Associates, the people who attended the stakeholder workshop in London, the respondents to the on-line survey and the internet forum participants.

EXECUTIVE SUMMARY

How does it help?

Licence-Exempt (LE) bands are areas of the spectrum that are very lightly regulated by Ofcom. Users do not need licences to transmit in these bands and this has led to diverse applications and services relying on these bands. Are the existing bands sufficient to meet the current demand for services? What kinds of problems are experienced in places where there is too much demand? Do all the different services co-exist easily or do they interfere with each other?

Understanding how the LE bands are being used is an important input to Ofcom's decision-making process when it comes to considering any proposed changes to spectrum allocation and management. The LE bands have been studied in previous years through measurement at fixed sites, computer modelling and sectoral studies. Those investigations indicated a continuing rise in usage, which raises questions about future growth and whether or not the regulatory approach to the LE bands will continue to be appropriate for the services in these bands.

One of the main services using LE spectrum is the wireless local area networking technology commonly referred to as WiFi. It is widely used as for data communications in laptop computers and handheld devices. This research concentrates on investigating the user experience of WiFi-based services as an indicator of the current state of an LE band that is playing an increasingly important role in the UK's communications infrastructure.

Explanation of the technology

This research concentrates on the 2.4 GHz band, which extends from 2.4 GHz to 2.483 GHz and is designated for Industrial, Scientific and Medical (ISM) use. It is a band that is used for many purposes, including:

- Wireless computer networking (WiFi, Bluetooth, ZigBee, mesh networks, etc.);
- Voice over Internet Protocol (VoIP) telephony;
- Gaming;
- Remote control;
- Audio Video (AV) senders and baby monitors.

The 2.4 GHz band also contains the band within which microwave ovens are allowed to operate. These devices are screened but some radio waves do still leak out and, whilst not at a level to be dangerous, they can cause interference to nearby communications.

WiFi networks

This research concentrates on the use of WiFi for computer networking. The term WiFi refers to a family of networking protocols, the first and simplest of which was IEEE 802.11b. More recent protocols, such as IEEE 802.11g and IEEE 802.11i, extend the original standard by allowing faster data rates, longer range, better multipath performance and improved security.

Normally WiFi communications are carried out via a central device called an Access Point (AP) in a mode of operation called *infrastructure mode*. Figure 1 illustrates infrastucture mode with a number of devices all communicating with the Internet and each other via the AP. It is normal practise to connect the AP to the Internet by a wired link, such as a domestic broadband connection.

It is also possible to configure WiFi networks in *ad hoc mode* which does not need an AP, but this mode is less common than infrastructure mode.

Figure 1 WiFi communications in infrastructure mode

WiFi frames

WiFi networks use the same Internet Protocol (IP) technologies as the rest of the Internet, but a dedicated layer of processing is added to allow communication by radio. From the point of view of this project it is this wireless communications layer that is of interest. This layer is called the *Link layer* or *Media Access Control* (MAC) layer.

Figure 2 shows the layers of communications involved in WiFi as a processing stack. The top three layers are those encountered in the wired Internet and are the application layer, transport layer and internet layer. The link layer is shown underneath these and, below that, there is the physical layer, which represents the radio transmitter/ receiver and the spectrum itself.

Internet	Application layer
	Transport layer
	Internet layer
Wireless data communications	Link (MAC) layer
2.4 GHz ISM band radio interface	Physical layer

Figure 2 WiFi protocol stack

Within the wired Internet, data is carried within the internet layer in short bursts called *packets*. A similar principle applies within the link layer, but the data bursts are called *frames*.

A frame is a single burst of data and all WiFi messages are transported between two devices by one or more frames. There are three main types of frames: management frames, control frames and data frames.

These frames can be readily recorded and analysed using commercially available test equipment and by computing devices that are running suitable software. There are actually over thirty different kinds of frames and it turns out to be useful to look at certain types of frame for measuring specific effects. In particular:

- **Beacon frames** are a type of management frame. They are transmitted at regular intervals by APs to advertise their presence to potential client devices. Counting the beacon frames gives an indication of how many APs are operating in an area.

- **User data frames** are a type of data frame. They carry the actual data generated by users. By counting the user data frames it is possible to estimate how much content is being passed across the wireless networks.

If a frame is lost for any reason, it can be retransmitted by the sending device. These retransmissions can be observed and used to estimate how many frames are being lost. This is one way of looking for problems in the WiFi networks.

Current state of the art and our research

This research is being performed in an environment in which the 2.4 GHz band is already being used heavily for WiFi and other services. Earlier survey work indicated the rise in usage of this band, so continued monitoring of this band will help to indicate when trouble spots are likely to appear and how bad they might be.

Previous surveys and modelling activities concentrated on obtaining estimates of the physical layer utilisation. This project takes a different approach and looks instead at the link layer.

A main thrust of this project has been to try to quantify the likely user experience in different areas. We have researched this concept by loading experimental networks and gauging the user experience at different levels of network traffic. From this investigation we have proposed a way of making measurements that can suggest the worst-case user experience in an area. The measurements can be made with a handheld device that is carried around an area of interest.

Frame rates

Counting the total number of WiFi frames received per second is one way of measuring how busy an area is. Considerable variation exists between the total frame rates at different sites, as shown in Figure 3, in which it can be seen that the total frame rate is higher in London than other places.

Furthermore, the proportions of the different frame types vary from site to site. In Figure 3 the frame types have been split into user data frames, beacon frames and 'other management and control' frames.

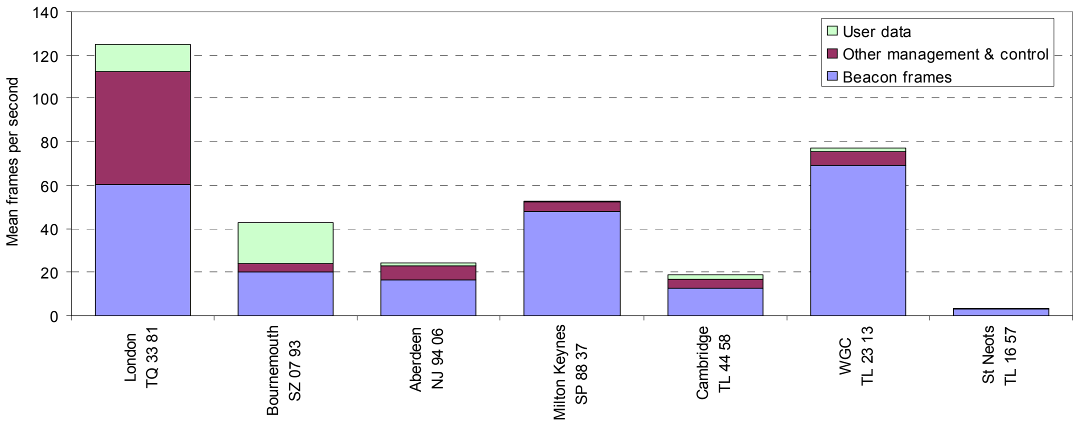

Figure 3 Proportions of frame types in different locations

We have found that it is rare for the user data frame rate to exceed 10% of the total frame rate. The implication is that over 90% of the frames are typically being used to keep the networks running and are not carrying user data. This is an important result, as it suggests that the existing WiFi protocols can use a substantial part of the 2.4 GHz band without carrying significant quantities of user-generated content.

Problems in the 2.4 GHz band

Gauging the levels of usage and network degradation have been the main thrusts of this research programme. These are two quite separate questions. A spectrum band can be very busy without users experiencing any problems. Conversely, a band can be lightly loaded but the users are not obtaining the service quality they expect.

It was clear from our research that the majority of problems experienced by WiFi network users are not spectrum-related. There are many ways in which the user experience can be affected, but only some of the problems can be attributed to working over a shared radio link. With 83 MHz of bandwidth available and a restricted range for WiFi services, the 2.4 GHz band can support many users simultaneously in an area. It is only in the centre of London that wireless congestion was found to be a significant problem.

With a very wide range of applications, technologies and services in the 2.4 GHz band, users will have correspondingly diverse expectations. This was confirmed by our investigations, which also revealed very strong beliefs and opinions held by many users.

Sometimes these beliefs can be supported by the evidence available, but this is not always the case. For example, there is a view that some domestic users generate excessive amounts of WiFi traffic, denying access to other users. Our research suggests that this is not the case, rather the affected parties are almost certainly seeing interference from non-WiFi devices such as microwave ovens, Audio Video (AV) senders, security cameras or baby monitors.

Estimating user experience

With so many different problems affecting users' experience of wireless networking in the 2.4 GHz band, producing a single number that can be meaningful to all users is impossible. However, given some assumptions, it is possible to give some measure of the worst case that people could expect in an area.

We define a parameter called the Mean Opinion Score Lower Bound (MOSLB) and calculate this number for individual 1 km grid squares. This can then be plotted on a map to indicate the areas where WiFi problems may be expected. An example is given in Figure 4 for the city of Cambridge. The worst areas are the two red squares, which is where we would expect most problems to be occuring.

Figure 5 shows the same area in 3D with the heights of each column indicating the average frame rate. It will be seen that the red columns are not the tallest ones. The areas where most problems are observed are not necessarily the busiest areas. This can be explained by interference. Areas where the WiFi services are experiencing a lot of interference are not necessarily those where there is a lot of WiFi usage.

Figure 4 MOSLB map of Cambridge

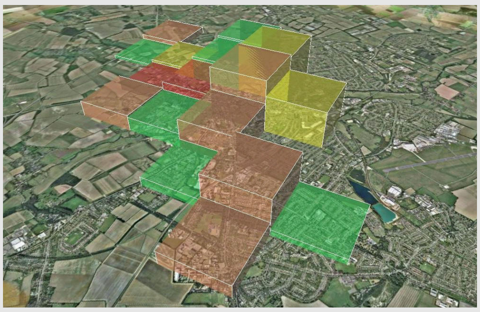

Figure 5 MOSLB and frame rate map of Cambridge

Interference and Congestion

Our measurements and experiments suggest that the WiFi networks are capable of carrying much more traffic than they actually do. However, interference between different types of devices is commonplace, leading to loss of service quality for many users.

Use of the word 'congestion' is misleading in this context, as it implies that performance degradation is due to high levels of usage, which is not supported by our measurements. The drop in performance in urban areas appears to be mainly the result of interference from other wireless devices, problems in the wired Internet and problems with device configuration.

These same effects were observed at different locations around the UK and we are led to the conclusion that interference between different kinds of radio in the 2.4 GHz band is commonplace. The greatest concentration of different radio types tends to occur in urban centres, so interference tends to increase with population density.

However, interference also occurs in low population density areas. It only requires a single device, such as an analogue video sender, to severely affect WiFi services within a short range, such that a single large building or cluster of houses can experience difficulties with using a single WiFi channel.

These interference problems are the indirect result of light regulation in the 2.4 GHz band. A plethora of radio types, which are not all designed via standardisation processes, means that peaceful co-existence does not arise organically. Co-existence must be enforced by some means if LE bands are to be shared effectively.

Conclusions

From a technical viewpoint the main conclusions from our research are:

- The majority of problems experienced by WiFi users are not spectrum-related. Users are likely to attribute their problems to congestion, but most of the time the problems are due to wired Internet problems or device configuration errors.

- Interference is commonplace and is a more important cause of wireless networking problems than congestion. In the long term this could be reduced by enforcing co-existence criteria via the standardisation committees. However, in the short term there are a lot of radio types in use and the interference problem is predicted to continue to increase.

- Inner city locations are extremely busy and do exhibit signs of congestion as well as interference. We expect this to be occuring in most large cities of the UK.

Interpreting the technical findings of the research in terms of the impact on the LE bands, we are led to the following views:

- The LE bands are an increasingly important part of the UK infrastructure. Organisations that use these bands should be aware that service quality cannot be guaranteed and that interference can appear suddenly when a nearby user turns on a different kind of radio device. The LE model means that both users may well suffer the effects of interference from the other and there is no legal basis for resolving conflicts.

- The lack of licence costs makes LE bands attractive to organisations such as local authorities for networking of public services in urban areas. Advice should be given on the need for resilience if LE bands are used for critical infrastructure services.

- There is a tendency for people to buy commercially-available equipment and adapt it to increase the range. Such modifications are likely to deny services to other users and should be discouraged. If the LE bands are to continue to be available to everyone then the power limits on devices need to be adhered to. Similarly, using standardised equipment should be encouraged to improve co-existence.

- Further work on communications standards should concentrate not on increasing data rate and range, but rather on:

    (a) Improving co-existence of different protocol families (which may require externally-funded, co-operative research between competing standardisation bodies and technology alliances);

    (b) Automatic configuration throughout the protocol stack to minimise user error and maximise the network efficiency and co-existence in shared LE bands;

    (c) Minimising the networking overhead so that the number of frames required for management and control is reduced;

    (d) Improving the ways in which networking protocols react to interference and congestion;

    (e) Removing conditions under which client devices can flood a network with, for example, null data frames.

- We see the co-existence issue as important enough to require a certification mechanism. A '2.4GHz friendly' logo attached to devices that adhere to co-existence standards would be an approach that manufacturers and consumers could understand and support. It would allow innovation by providing a target to achieve for new devices. The baby monitor market is one example where this approach has already arisen through market forces, with 'WiFi-friendly' appearing on the datasheets for the new generation of digital video monitors. A '2.4 GHz friendly' logo would tap into such market pressures and lead to improved co-existence of all types of radio in this band.

CHANGE HISTORY

Version	Date	IR	Comments
1	2/2/09	-	First issue for Ofcom approval
2	27/3/09	-	Second issue
3	1/4/09	-	Third issue

This report was commissioned by Ofcom to provide an independent view on issues relevant to its duties as regulator for the UK communication industry, for example issues of future technology or efficient use of the radio spectrum in the United Kingdom. The assumptions, conclusions and recommendations expressed in this report are entirely those of the contractors and should not be attributed to Ofcom.

DISTRIBUTION

Copy No 1	William Webb	Ofcom Riverside House 2a Southwark Bridge Road London SE1 9HA
Copy No 2	Gary Clemo	Ofcom Riverside House 2a Southwark Bridge Road London SE1 9HA
Copy No 3	Cesar Gutierrez	Ofcom Riverside House 2a Southwark Bridge Road London SE1 9HA
Copy No 4	Project File	Mass Consultants Limited

CONTENTS

1 INTRODUCTION

This is the final report on the project entitled 'Estimating the Utilisation of Key Licence-Exempt Spectrum Bands' which has been undertaken for Ofcom by Mass Consultants Ltd. (MASS).

The project aimed to measure congestion in the 2.4 GHz Industrial, Scientific & Medical (ISM) band by concentrating on the performance of the IEEE 802.11 'WiFi' networks that have become commonplace across the UK. It canvassed the opinions of various stakeholders, made measurements at various locations and performed laboratory experiments to investigate the state of the 2.4 GHz band.

This report is structured as follows: Section 2 gives background material relevant to the project; Section 3 presents the general concepts behind measuring utilisation and looking for network problems; Section 4 looks at the various factors that can affect user experience; Section 5 covers the monitoring survey and Section 6 covers the various experiments carried out in the laboratory. Section 7 presents the analysis of the results from a purely technical perspective then Section 8 introduces our proposal for relating monitoring statistics to user experience. Section 9 identifies some significant trends for the future of this LE band. Finally, our conclusions are presented in Section 10. Appendix A gives the results of the online questionnaire and Appendix B considers how frame rates should be estimated. Appendix C presents the results of the field survey as a series of colour-coded maps of the areas visited.

The remainder of this introduction summarises the research question and the key findings.

1.1 RESEARCH QUESTION

Ofcom relies on understanding of the current state of the spectrum in order to make policy decisions on future regulation. With lower power, shorter range devices it becomes increasingly difficult to obtain a clear, up-to-date picture and monitoring systems are necessary to provide this information.

Using numerous, relatively inexpensive, handheld receivers offers the potential of supplying the necessary survey data in real-time with good geographical coverage. This project therefore set out to investigate the validity of such an approach.

The following research question was posed at the start of the project:

Is it possible to define one or more technical measures of network congestion that can be obtained by passive monitoring and can be easily related to user experience of network congestion?

1.2 RESEARCH METHOD

The research question has a number of elements to it. Passive monitoring using conventional equipment is well-established (Section 2.3). An alternative was adopted for this project using small, handheld receivers to do the monitoring. It was necessary to consider a range of technical measurements available from the receivers and evaluate the measurements for the task of estimating user experience.

These various problems were dealt with in a variety of ways:

1 Stakeholder opinions were sought via meetings, a questionnaire and internet forums. These enquiries revealed the ways in which users perceived network degradation problems;

2 A small number of handheld receivers were used to gather data in UK towns and cities. This exercise enabled us to evaluate the logistics required to obtain good quality results;

3 Laboratory experiments were performed to better understand the behaviour of heterogeneous WiFi networks in various conditions. These experiments allowed the user experience to be assessed with a network in different states of network degradation which could not be carried out in the real world;

4 Different metrics were evaluated for their potential to be used as indicators of user experience. A single metric was selected as the recommended way of categorising different environments.

1.3 KEY FINDINGS

1	**Users experience many problems with WiFi networking, but most of these are due to problems other than wireless network degradation;**
2	**Where there is wireless network degradation this is usually due to interference rather than congestion;**
3	**Large city centres are a special case where both interference and congestion are occurring. This is not seen in smaller towns and cities.**

After considering the published literature, experimental results, survey data and stakeholder opinions we have concluded that there is no simple relationship between passive measurements and the user experience that can be expected. There are simply too many variables including, but not limited to: effective wireless range, wired and wireless network protocols, performance of Internet resources, behaviour of application software and user expectations.

We have, however, proposed that a lower bound on user experience can be estimated and that this is an effective means of identifying areas where there may be problems in the Licence-Exempt (LE) bands. This metric is based on the concept of relating network performance degradation to a Mean Opinion Score (MOS), which is a five point rating scale already used in the audio and video industries. We refer to this metric as the MOS Lower Bound (MOSLB).

Case Study - London Railway Stations

Liverpool Street is one of the main railway stations in central London and falls within the City of London district. Macdonald (2004) reported a good connection to an access point here, but Cochrane (2007) carried out a survey of the area where he found a very high network density and reported that connections at the station are difficult to establish.

As part of our survey we visited the station and the area around it twice. Attempts were made to connect to various access points but without success.

Measurements of the WiFi traffic in the Liverpool Street area showed that it is one of the busiest of the locations surveyed. In grid square TQ3381 we measured a network density of 657 BSSIDs/ch/km^2. Similar locations were Marylebone (TQ3080) with a network density of 733 BSSIDs/ch/km^2 and Euston/Kings Cross (TQ2982) which had 749 BSSIDs/ch/km^2. The London railway stations typically had 50% more networks per square kilometre than any other location.

Furthermore we have found that the common concept of congestion is inappropriate as it implies deterioration in performance with increasing load offered to the networks. Our results suggest that, whilst congestion can occur in a heterogeneous WiFi network, this happens at a level of activity significantly in excess of that seen in practice. There are, however, many sources of interference that reduce the ability of such networks to provide satisfactory levels of service quality and these are only weakly correlated with the numbers of users and access points in an area. We prefer to avoid using the word 'congestion' and use 'degradation' instead, as this does not imply a causal link between a problem and the level of network activity.

The taxonomy of section 4.2 categorises different kinds of network degradation to explain how diverse the problems are. Definitions of the terms congestion, interference and degradation, as used in this report, are given in section 12.

Case Study - Audio Video Senders

2.4 GHz AV senders, wireless security cameras and baby video monitors are readily available in the UK and are inexpensive. They can be used by anyone who wants to relay a TV signal over a short distance without installing wiring.

These devices typically use FM modulation and transmit at a power level of up to 10 mW. WiFi devices transmit at up to 100 mW, but spread that power over a wider band.

Our experiments show that AV senders cause serious disruption to WiFi services. At the same time, WiFi signals interfere with the AV signals.

The effect of AV senders on WiFi clients is distinctive and repeatable. As these devices are inexpensive and readily available in supermarkets, we conclude that the majority of the accusations levelled at residential users for abusing the WiFi networks are actually caused by the installation of AV senders or similar devices. Such devices can easily deprive neighbours of use of a WiFi channel and, at the same time, give the impression of overuse of the channel.

We have found that the MOSLB metric falls to an unacceptable level (i.e. level 1) in the centre of London. It is likely that the MOSLB is also very low in most city centres, although more survey data is needed to confirm this prediction.

The MOSLB also falls to an unacceptable level at isolated sites that may, or may not, be densely populated. It is highly likely that such sites are being affected by sources of interference, such as audio video (AV) senders, wireless security cameras, baby monitors and microwave ovens. We have seen this type of effect in the vicinity of a number of restaurants which suggests that microwave ovens may be to blame in those cases.

2 BACKGROUND

This section looks at three very different viewpoints which contribute to understanding of Licence-Exempt (LE) band utilisation:

- The development of new LE band protocols and services by scientists and engineers is well documented via technical journals, books and international standards. Section 2.1 therefore looks at the available literature;
- Anecdotal evidence of problems comes from users and other stakeholders. Section 2.2 contains the results of our investigations into various people's anecdotes, beliefs and opinions;
- MASS has itself carried out spectrum activity monitoring in these bands on behalf of the Radiocommunications Agency and Ofcom. Section 2.3 summarises this work.

2.1 LITERATURE SEARCH

There is a wide body of literature concerning interoperability of IEEE protocols, interference in LE bands and other network degradation issues.

An important distinction between the academic papers is in the environments used. Our particular interest is in the urban environment. Some research has already been done in such areas, but other researchers have concentrated on other types of environment such as university campuses, conferences or propagation test ranges. A lot of the research relies on theoretical models, both analytical and simulated.

2.1.1 Urban environments

Leeson *et al* (2000) looked at the future use of the 2.4 GHz ISM band and made a number of conclusions, including that there was a potential for interference between IEEE 802.11b, Electronic News Gathering and Outside Broadcast (ENG/OB) equipment and Radio Fixed Access (RFA) systems. They predicted a density of 40 AP/km^2 by 2005 and said that there would be problems if the density of outdoor networks exceeded 1 AP/km^2.

A flurry of interest in WiFi problems occurred around the period in 2003 when Kastner, cited in Sutherland (2003), published a report provocatively entitled "The urban WiFi crash of 2004". This predicted major problems with WiFi networks in 2004. The report was widely cited on the web at the time (Gardner, 2003) (Gubbins, 2003) (Sutherland, 2003) and authors such as Kewney (2003) were reporting similar concerns. The Kastner report created headlines but was not universally accepted. Molta (2003), for example, disagreed with the technical assumptions made in the Kastner report.

Brik, *et al* (2008) recently carried out extensive monitoring of a large-scale, urban mesh network that used IEEE 802.11b/g as its access method. The conclusions of that study were that the backbone network performed significantly better than the access network and that the disparity in performance was “mostly a result of unmitigated interference in 2.4 GHz spectrum in urban settings.”

The results of our survey and laboratory experiments suggest that, in the absence of interference, WiFi networks would support much higher throughput than they currently do. The current situation in central London does, however, seem to support both Kastner's predictions and Brik's measurements, with the potential for poor service quality being widespread. Hence we conclude that, whilst the "urban crash" may not have occurred with quite the drama expected in 2004, central London (and maybe other large urban centres) is now experiencing the problems foreseen in 2003. These are largely due, not to the density of the networks, but to other problems, such as interference from non-WiFi sources. Congestion within WiFi networks appears momentarily and, in some areas such as Liverpool Street station, is probably a persistent problem.

Cunningham and Grout (2007) have performed passive monitoring with an emphasis on determining how many networks were unsecured. They also considered the issue of determining the best route to follow to cover the area at lowest cost, which is an important consideration when planning a field survey.

2.1.2 Conferences

Large conferences appear to be a specific case where WiFi congestion has been a significant issue. In recent years there has been a growing trend for delegates to use their laptops extensively during large conferences and the organisers have to deploy sufficient internet access bandwidth to satisfy the demands.

The following are good examples of papers based on measurements made at large conferences:

Rodrig *et al* (2005)	SIGCOMM conference, 2004
Jardosh *et al* (2005a, 2005b)	Internet Engineering Task Force meeting, 2005
Raghavendra *et al* (2005)	Internet Engineering Task Force meeting, 2006

Rodrig *et al* (2005) found that the overheads of the WiFi protocols are high. In their studies they found that only 40% of the data transmitted was user data, the remainder was used for control and management purposes. They also found that client devices change their data rates very frequently and that retransmissions are common.

Jardosh *et al* (2005a, 2005b) carried out measurements at a large conference and developed monitoring methods that gave good quality measurements of network performance. In their case, however, they had knowledge of the network infrastructure, which is not the case in the urban environment. Their measure of the link reliability was based on the number of missing beacon frames, but this method was not practical in our application. The same authors made relevant observations on the use of RTS/CTS and rate adaptation (Section 4.2.2).

Raghavendra *et al* (2005) concentrated on the handoff events that occur in a heavily loaded conference environment. They found that, even when the clients are largely static, the IEEE 802.11 protocols tend to cause a high rate of handoff events which are incorrectly initiated and add to the network traffic. These erroneous handoffs therefore exacerbate any congestion already present. Each reassociation following a handoff, in addition to generating unnecessary network traffic, also impacts the application layer, so the user is directly affected by such events.

2.1.3 University campus and laboratory studies

Heusse *et al* (2003) set up an experimental network and demonstrated that the throughput is much lower than the nominal bit rate. They also showed that rate adaptation tends to aggravate congestion.

Hansell *et* al (2004) used an experimental network to investigate throughput characteristics of one WiFi network in the presence of other WiFi and Bluetooth networks.

Yeo *et al* (2005) used passive monitoring in a computer science department and demonstrated that it can produce useful statistics despite concerns over packet loss and lack of knowledge of the transmitted signals. They point out that the passive monitoring approach does not require network sniffers to be attached to the wired networks and can produce detailed information on wireless-side statistics.

Franceschinis *et al* (2005) ran tests on a WiFi network at the electronics department in the Politecnico di Torino, Italy. They showed that TCP connections were sensitive to the settings of the maximum TCP congestion window. These researchers disagreed with all other results concerning RTS/CTS and decided that use of the protocol improved TCP throughput by reducing the number of frame collisions. They also concluded that unfairness could be different between uplink and downlinks.

The study of wireless networks on university campuses has the advantage for the researchers that the infrastructure can be completely understood. Mahanti *et al* (2007) used a passive measurement method over a period of six weeks. This length of time indicates another advantage of campus studies where it is relatively inexpensive to gather long term statistics.

The Mahanti paper analyses the network traffic with a view to understanding social behaviour. The data they collected was used to work out favourite web sites and preferred locations in the campus for internet use. Within an organisation it is possible to obtain consent from all participants for such analyses, but this is not practical in an urban environment.

Kemerlis *et al* (2006) reported on measurements made on a test network at Athens University. Their results showed that, when TCP is being used, the clients with the highest signal strength get unfair access to the wireless channel. This effect is attributed to the increased rate of frame loss on lower strength links, which TCP interprets as being congested. These researchers also reaffirmed the findings that RTS/CTS reduces the available throughput.

2.1.4 Propagation studies

Anastasi *et al* (2004) studied various aspects of IEEE 802.11 network propagation including the hidden node problem, for which they concluded that the RTS/CTS protocol was not an adequate solution.

Barsocchi *et al* (2005) presented a method for performing WiFi channel measurements at the link layer. Their approach allowed the measurement of received power level, frame loss and packet delay. The results compared well with a two ray propagation model in a rural environment.

2.1.5 Theoretical studies

Jun *et al* (2003) calculated the theoretical maximum throughputs of the IEEE 802.11 protocols. They concluded that the achievable throughput is typically much less than the 'headline rate' and reduces for smaller frame sizes. They also concluded that using RTS/CTS reduces the bandwidth efficiency.

Battiti *et al* (2003) proposed a method for controlling access to Access Points based on a pricing system, which they called Price-based Congestion Control. It was recognised that users may complain if their Quality of Service varied depending on the number of users attached to an AP. The object of the proposed algorithm was to maintain the Quality of Service for users who were prepared to pay for it.

Chatzimisios *et al* (2005) used a Markov chain model to analyse the performance of the link layer performance, with and without RTS/CTS. This allowed them to make recommendations on the size of contention window and retry limit. The implications of this and similar work by other researchers is that the link layer performance of these networks is very dependent on the settings of various parameters. In a heterogeneous, densely-populated environment it is highly unlikely that all the network devices are configured for optimum performance.

Golmie (SA) considered the interference between Bluetooth and WiFi and summarised the ways of improving coexistence between them. Recommendations were made on protocol parameter settings based on the results of simulations.

Awoniyi and Tobagi (2004) modelled the IEEE 802.11a channel with a view to understanding the impact of fading on VoIP services and concluded that fading has a significant impact. They relate their results to a Mean Opinion Score (MOS) which is a concept extended to general browsing in section 8 of this report.

Hansell *et al* (2004) simulated multiple AP scenarios using the 'N-Systems' method and concluded that the network density could reach 24.79 AP/km^2 before congestion became a significant issue. Their simulations showed a marked reduction in this figure if Bluetooth networks were introduced as interferers.

2.1.6 Summary

The extensive literature on the development of the WiFi protocol and related areas suggests that some facets of the logic employed are flawed when interference and congestion occur.

Whenever an error occurs the protocols tend to assume that the signal strength is low. The strategies for dealing with interference and congestion should be different from those for dealing with low signal strength. The logic currently used in many devices tends to make congestion worse rather than better.

Interference is a complex subject and has only been studied in detail for certain combinations of protocol, e.g. Bluetooth versus WiFi. Generally the assumption seems to be made that other sources of interference can be handled as noise sources, which is unlikely to be the case when interferers are communications signals.

2.2 ANECDOTAL EVIDENCE FOR PROBLEMS

A number of different approaches have been followed in order to see what anecdotal evidence exists for problems in the 2.4 GHz ISM band. The approach taken has been to try to stimulate discussions and then record people's opinions. In this section the results of the investigations are presented.

The following approaches have been adopted during this phase of the project:

1 Literature search;

2 Forums;

3 On-line survey;

4 Stakeholder discussions.

The overall impression is that there are a number of widely-held beliefs and a number of areas of debate. The opinions expressed by individuals are frequently made with the benefit of practical experience but are not backed up by documentary evidence.

2.2.1 Forums

Questions about congestion in the 2.4 GHz ISM band were posted on a number of web forums. Eleven forums were posted to and useful threads appeared on five of these.

The main observation arising from this investigation is that there are several widely-held beliefs, the main ones being:

1 Most urban areas are suffering from congestion;

2 Most commercial networks are properly set up;

3 Residential users are the cause of congestion;

4 Residential users do not change the default channel;

5 Residential and non-residential users should be on different channels;

6 More channels should be added.

These beliefs are sometimes expressed quite forcefully, although there is little concrete evidence to support the assertions. They appear to be regarded as self-evident truths by some forum participants.

In addition to the above beliefs, there are a number of debates with people arguing for and against each of the following:

1 WiFi access should be free at the point of use;

2 Residential users should use wired access;

3 Intelligent channel selection within routers is needed to avoid congestion.

Our responses to these beliefs and debates are summarised in the tables below:

Belief	Our response
1. Most urban areas are suffering from degradation	Our results tend to support this belief. Extensive WiFi problems have been observed in the centre of London and it is likely that similar situations occur in other large cities.
2. Most commercial networks are properly set up	This has not been investigated.
3. Residential users are the cause of congestion	We do not support this belief, as we have found no evidence to suggest that there are problems due to overuse by residential users. We do, however, suspect that such opinions arise from cases of interference in residential areas, with AV senders being the main cause of network degradation (Section 7.6.1).
4. Residential users do not change the default channel	This may be the case, but our measurements of channel balance (Section 8.4) indicate that this is currently not a major problem in the UK.
5. Residential and non-residential users should be on different channels	This belief is driven by an assertion that residential users are abusing WiFi networks, but our results do not support this assertion. We are concerned that organisations may be deploying mission-critical or safety-critical networks in LE bands, where there is little regulatory protection from interference and protocol incompatibility.

Belief	Our response
6. More channels should be added	This belief arises from users who are experiencing lack of service quality in the LE bands. Our results suggest that there is sufficient bandwidth in the 2.4 GHz band, but that much of it is being lost through interference of various kinds. Adding more LE channels at lower frequencies is unlikely to produce a solution for wide bandwidth applications, as the LE spectrum tends to be inefficiently used. Adding more channels at higher frequencies would be useful providing that the power limits are set such that interference between neighbouring networks is minimised. As the 5 GHz LE band is already available, our recommendation is to use that for low power, high bandwidth applications.

Debate	Our response
1. WiFi access should be free at the point of use	This debate does not relate directly to the issue of network degradation, but the free availability of access may lead to increased demand and therefore the increased likelihood of congestion.
2. Residential users should use wired access	This is a corollary of the belief that residential users are the main cause of congestion. The proponents assert that they should not, therefore, be allowed to use WiFi networks. Our results do not support this argument.
3. Intelligent channel selection within routers is needed to avoid congestion	This debate arises from the belief that most users do not change the default channel when setting up their WiFi networks. Our measurements of channel balance (Section 8.4) suggest that this is not a major problem in the UK. Automatic channel selection is already implemented in many devices and will become more widespread in the future

2.2.2 On-line survey

An on-line survey was sent to a range of people within industry, academia and the general public. A total of 54 people completed the questionnaire, the results of which are given in detail in Appendix A. Whilst only a small sample, this questionnaire did produce valuable insights into the stakeholders' opinions.

The opinions expressed by individual respondents largely reflected those found in the forums and these have been addressed in Section 2.2.1. In addition to those beliefs and debates, the following questions arose:

1. Should the power limit be raised, kept the same or lowered?

   Tension is ongoing between the needs of different applications for WiFi services. Individuals often want to increase the range of their connections (Goodwin, 2007) (Freenet, 2008), but, at the same time, they want other people's devices to use lower power to minimise any adverse effects on their own connections.

   67% of the questionnaire respondents thought that the current power limits should not be removed. Engineers and regulators understand the issues and know that a compromise position has to be enforced so that everyone can benefit from the LE spectrum.

2. Would moving to 5GHz be beneficial for some users?

   This area has not been investigated in detail, but it appears that some people think that 5 GHz would offer advantages for some applications, although it may only be a short-term solution. We expect that this will start to become more of an issue if the 5 GHz band starts to be used for high definition television via the Wireless High Definition Interface (WHDI) standard.

   The respondents pointed out that, at 5 GHz, applications will be restricted to short-range and line-of-sight, which is restrictive compared to 2.4 GHz. It implies the need for more Access Points (AP) which implies increased infrastructure cost. A 5 GHz network will therefore be more suitable for business use than domestic use.

   The technical feasibility of deploying a greater density of APs at 5 GHz will depend on the local propagation environment, equipment battery life and other factors which are beyond the scope of this study.

2.2.3 Stakeholder discussions

A number of stakeholders have been approached in order to canvass their opinions. Notably, the following organisations have expressed their viewpoints:

- Intellect (meeting)
- Open Spectrum (email)
- Cochrane Associates (email)

The culmination of these discussions was a workshop held at the Ofcom headquarters in London on 19th September 2008 to which a number of stakeholders were invited.

The following paragraphs summarise the main outcomes of these discussions.

2.2.3.1 Meeting at Intellect

> "Intellect is the trade association for the UK technology industry."
>
> "Intellect provides a collective voice for its members and drives connections with government and business to create a commercial environment in which they can thrive. Intellect represents over 800 companies ranging from SMEs to multinationals. "
>
> (www.intellectuk.org/content/view/23/3, accessed 10 Sep 2008

A meeting was held at the Intellect offices in London on 24th June 2008. After presenting the current project's objectives a lengthy discussion followed which raised many issues, of which the following is a summary:

- Applications such as the BBC iPlayer are predicted to cause the downfall of WiFi, as excessive use of video services will lead to unacceptable levels of congestion in urban areas. "Up to half a million programmes a day are being streamed or downloaded through the BBC iPlayer" according to informitv (2008). Evans (2008) wrote that "ISPs are unhappy for the simple reason that usage is starting to clog their networks...". With many home users now using laptops to connect wirelessly, it is reasonable to assume that an increased use of internet video services will result in increased data on domestic WiFi networks.
- A report on convergence was compiled by Intellect and sent to the Convergence Think Tank in April 2008, highlighting the problem of internet congestion (Intellect, 2008);

- Intellect believes that the default channel problem is significant. Most users do not know how to change the channel and would not understand the ramifications of doing so. Equipment manufacturers could improve the configuration software and user information they provide to improve this situation;

    - After analysing the data from the survey, our results do not support this belief (see section 8.4) and suggest that the channels are reasonably balanced in the busier areas.

- In a similar vein, Intellect believes that many users are not aware of the need to secure their networks and need to be prompted, in a meaningful way, to set up the encryption mechanisms available on all routers;

- Intellect believes that network layer congestion and ISP-related problems are currently more significant than physical layer congestion;

    - Our investigations support this. We propose (section 4.2) that the term 'degradation' is more appropriate than 'congestion' when talking about the wireless physical and link layers and that there are many more problems that users encounter when using WiFi services.

- Intellect stated that the use of proprietary signal formats is unhelpful as these can interfere with the standard modulations and protocols. AV senders are seen as a significant interferer as they are bandwidth-inefficient;

    - Again we support this belief. In the most extreme cases, AV senders are not polite at all and can completely deny access to a WiFi channel (section 7.6.1). Other devices, such as wireless security cameras and baby monitors use similar technologies to AV senders and are likely to have similar effects on WiFi.

- The London borough of Islington was thought to be extremely busy and a good place to look for congestion;

    - We have not visited Islington, but the central London results suggest that Islington would have high levels of network degradation.

- A specific issue relating to polite protocols was raised. Polite protocols are a problem on inter-access point links in trunked systems as one generally wants such links to run at maximum capacity. This issue is largely outside the scope of the current project and has been discussed previously with Ofcom. It has also been a subject of an Ofcom research programme (Richardson, 2006) (Ofcom, 2007).

2.2.3.2 Email discussion with Open Spectrum

> "Open Spectrum UK seeks to engage the public interest agenda for the exploitation of the strategic national resource of the radio spectrum. "
>
> "Open Spectrum UK argues for a balance of the commercial and the public interest in access to and use of the radio spectrum. "
>
> "Open Spectrum UK plans a programme of activities to engage informed debate upon the future of spectrum access and communications rights in the UK."
>
> (http://openspectrum.org.uk/wiki/wikka.php?wakka=OpenSpectrumUKAims, accessed 16 Sep 2008)

Discussions have taken place by email with Open Spectrum. The salient points are as follows:

- Horvitz (2008) presents the case for the 2.4 GHz ISM band becoming "super-saturated". It refers to work done by Aegis and Transfinite, Peter Cochrane, RSA and MASS. He makes the point that it is difficult to resolve the measurements and predictions of these sources.
  - These studies were discussed at the stakeholder workshop (section 2.2.3.4) and the differences between the results were explained.
- Horvitz (2008) suggested the use of remote controlled miniature helicopters to survey urban areas.
  - This is similar to our approach of using small, handheld receivers (section 5.1). It would be very difficult to use miniature helicopters within buildings, so we prefer the method used in our project.
- Spectrum occupancy monitoring has been carried out by Aachen University using equipment very similar to that used on previous surveys for Ofcom (Wellens and Mähönen, 2007). Their results suggest spectrum occupancies approaching 100% in outdoor measurements between 20MHz and 3 GHz.
  - The utilisation levels are very dependent on location, especially at higher frequencies. The Aachen results are not at odds with our findings in this or earlier studies.
- Bob Horvitz supplied a copy of a thesis which contains additional information on the usage of the 2.4 GHz ISM band. In his PhD. thesis Petrin (2005) proposed the use of the 'white space' in this band to maximise its use.
  - Our results suggest that, outside the major cities, there will generally be plenty of spare bandwidth for smart radio systems to exploit. In the busiest areas, however, it is unlikely that a very polite protocol would be able to find white space in which to operate.

2.2.3.3 Email discussion with Cochrane Associates

Peter Cochrane has written about WiFi availability and usage in various places, including on his blog (www.cochrane.org.uk/opinion). A number of salient points came out of an email discussion on the subject:

- With all digital communications systems it is difficult to discern the difference between interference, congestion, bad engineering and poor site location. You need the appropriate test equipment to resolve these issues;
- Peter has worked at locations such as the Institute of Directors in London with over twenty laptops in one room without any problems. Other locations, such as Liverpool Street station and Heathrow airport lounges are always problematic;
- 'Shock excitement' of receivers due to nearby transmitters could be a significant issue. This has not been investigated in this project but may be of interest for future research. The problem relates to the dead time that receivers can experience when recovering from a particularly strong signal.

2.2.3.4 Stakeholder Workshop

A workshop was held at Ofcom's Riverside House offices in London on 19th September 2008. At this meeting a number of presentations were given by MASS describing the project and several discussion sessions were held.

The workshop was attended by representatives from:

- Barclay Associates
- BT
- David Hall Systems
- Huber + Suhner (UK)
- Intellect
- MASS
- MLL Telecom
- Motorola
- Nokia
- Ofcom
- Open Spectrum
- Radio Society of Great Britain
- T-Mobile

Many points were raised during this meeting. The following is a list of those that were particularly pertinent to this project:

1 MASS confirmed that the survey covered both indoor and outdoor measurements. In general the outdoor environment in London appears busier than indoors;

2 The issue of data rate adaptation was raised (see Section 4.2.2) and whether this was a justification for using cognitive radio. It was agreed, however, that relatively straightforward changes to the existing protocol would be sufficient to address the problem without requiring the complexity of a cognitive radio solution;

3 The idea of mandating the use of Listen Before Talk (LBT) in the 2.4 GHz band was discussed. It was generally felt that the WiFi protocols are already 'polite' and that the problems were caused more by other services that did not use any form of politeness. Audio video senders, for example, transmit continuously without any regard to other users of the band;

4 The MOS concept (see Section 8.1) was described and was generally well-received, although there was concern that passive measurements would not be able to produce enough information for a MOS to be calculated (e.g. latency has a big impact on voice communications and cannot be measured passively). This was at a stage in the project when MASS was trying to estimate the MOS directly and before the decision was made to use the MOSLB instead (see Section 8.2);

5 Attendees were interested in looking at the 5 GHz band and other communications protocols in the 2.4 GHz Band (e.g. Bluetooth). Also, studying the use of the LE bands for mission-critical and safety-critical services would be of interest;

6 The issue of regulation in the LE bands was raised. From a regulatory point of view the main question is whether more spectrum needs to be released or the existing spectrum needs to be managed more proactively.

2.2.4 Summary

The opinions of various stakeholders and members of the public have been sought with a view to understanding the main issues currently being encountered in the 2.4 GHz band.

It is clear that problems are occurring, but the causes of these problems are not always well understood. This is because the communications devices do not give accurate explanations of the problems. People will therefore look at the evidence available and blame their neighbours for being greedy or point to the presence of too many network nodes in the area.

There is a wide spread of public opinion. Perhaps surprising was the finding that there are some strongly-held beliefs about what is going wrong when WiFi services don't deliver the desired performance.

Our investigations have supported some of these beliefs but not others. We do not, for example, support the belief that domestic users, through excessive usage, are to blame for the majority of problems. On the other hand, our results do support the belief that degradation in the large cities exists.

Similarly, there is a widespread view that users do not change the default channel on their domestic networks, which leads to excessive use of one channel over the others. Again our results do not support this belief, with good channel balance seen at the busier sites.

The increased use of the 2.4 GHz band is attributable to a growth in the different types of service available and also to the increased distribution of video content. When proprietary signal formats are used, which are not designed to be compatible with other signals, then there is potential for interference and this is supported by the reports of the people consulted.

In support of the findings from the literature search, it was found that people believed that rate adaptation is a problem in congested environments. The implication is that further work is needed in the development of the networking protocols to handle congestion more effectively.

The idea of using a MOS scale to indicate user experience was well-received. This encouraged MASS to investigate the approach further.

2.3 LICENCE-EXEMPT BAND MONITORING

MASS has built measurement systems and carried out field measurements in LE bands as part of the following projects for the Radiocommunications Agency and Ofcom:

- 2.4GHz Measurement Study (2003)
- Autonomous Interference Monitoring System (AIMS) phase 2 (2006-7)

The results of these studies were published on the Ofcom website and are available in Day and Wagstaff (2003) and Wagstaff and Merricks (2007) respectively. Both projects performed spectrum utilisation measurement in the physical layer.

2.3.1 2.4 GHz measurement study (2003)

The survey in 2003 looked only at the 2.4GHz ISM band. Eleven static sites were monitored and two roving surveys were performed. Figure 6 is taken from Day and Wagstaff (2003) and shows the sites monitored around the Cambridge area.

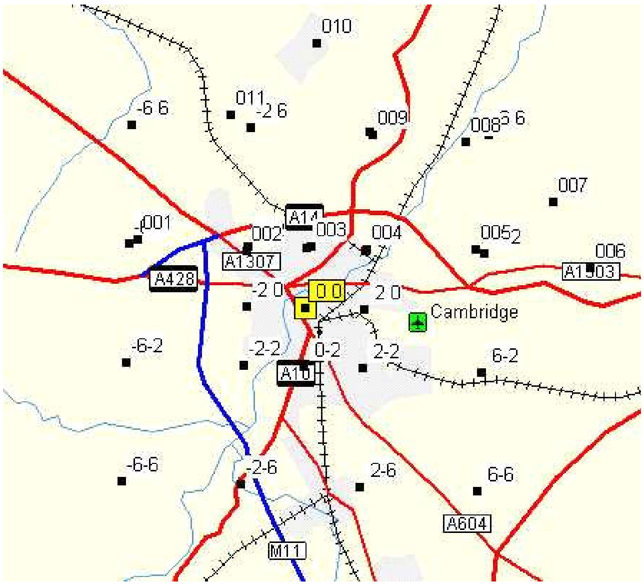

Figure 6 Large-scale roving survey

The sites surveyed were varied and included sites where 2.4GHz usage was expected to be high, including Heathrow airport (Figure 7). Considerable variation in spectrum usage was recorded between the sites. WiFi usage was, at that time, only just starting to be observed and microwave oven leakage was found to be the most prevalent signal in the environment.

Figure 7 Heathrow airport monitoring locations

This study included the analysis of different devices that radiate in the 2.4GHz band including WiFi, Bluetooth, video senders, motion detectors and microwave ovens. Figure 8 is taken from Day and Wagstaff (2003) and shows just one example of the statistics captured for a multiple Wireless Local Area Network (WLAN) scenario.

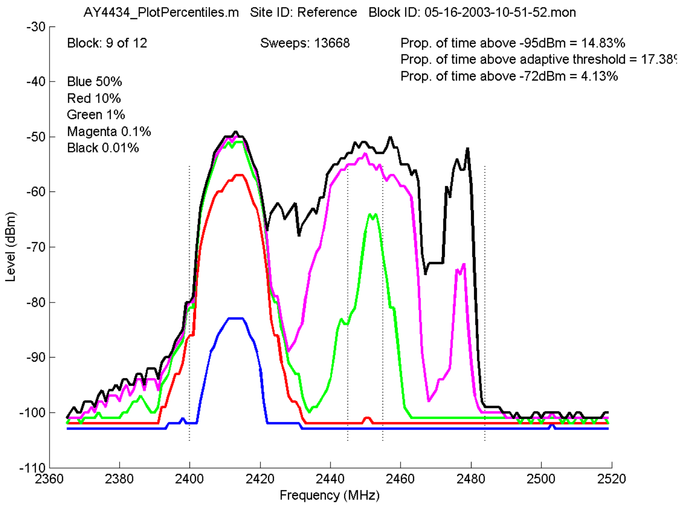

Figure 8 Multiple WLANs

2.3.2 AIMS phase 2 (2006-2007)

An LE band measurement exercise was included as part of the AIMS phase 2 programme. This involved making measurements in all seventeen LE bands between 100MHz and 10.6 GHz at 20 different sites. Figure 9 shows one of the sites at which the LE bands were surveyed.

Figure 9 AIMS monitoring LE bands at Sheffield University

Figure 10 is taken from Volume 1 of Wagstaff and Merricks (2007) and shows the range of average utilisations of the LE bands surveyed in 2006.

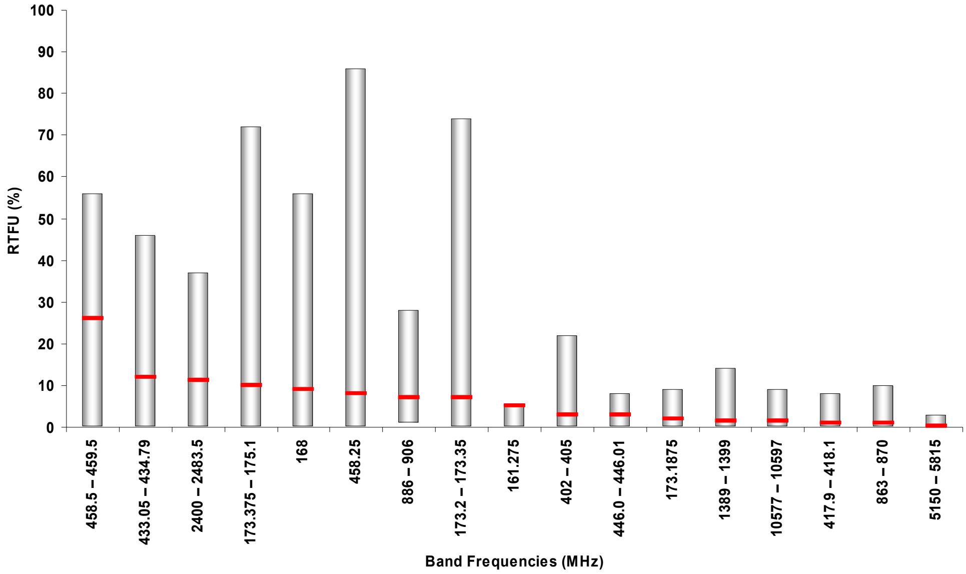

Figure 10 LE Band Relative Utilisation – Sorted by Mean Relative Utilisation

The main finding was that the LE bands were not, on average, heavily used. It is, however, important to stress that peak utilisation was much higher than the average in some cases.

The band displaying the highest usage was the 458.5 – 459.5 MHz band, designated for use by telemetry systems. Use of RF heating equipment in the 886 – 906 MHz band was also widespread, but the low duty cycle meant that utilisation appeared low.

As far as the 2.4 GHz ISM band was concerned, our surveys in 2003 and 2006 allowed us to compare utilisation at the three (indoor) sites visited in both cases. We found that the average utilisation was 5.7% in 2003 rising to 14.3% in 2006. Despite these increases, measurements in the 2.4 GHz ISM band only showed heavy usage at one of these three sites.

With the background of these measurements of utilisation in the physical layer, we now turn to defining and measuring network degradation. Section 3 describes how it is necessary to make measurements above the physical layer to supplement the information available from utilisation measurements. Section 4 presents an analysis of the types of network degradation that may be considered when using higher layer measurements.

3 APPROACHES TO MEASURING UTILISATION AND DEGRADATION

This section looks at the overall philosophy of carrying out field surveys to measure wireless network performance. It is shown that network performance is best viewed in terms of the two quantities: utilisation and degradation. If the utilisation measurements are sufficient then measurements can be made at the physical layer. If, however, the degradation also needs to be assessed, then the measurements need to be made at the link layer.

3.1 PHYSICAL LAYER MONITORING

Conventional spectrum monitoring, as described in Section 2.3, concentrates on measuring the amount of utilisation at the physical layer. Systems such as the Unattended Monitoring System (UMS) and Autonomous Interference Monitoring System (AIMS) measure the amount of time that the power in a spectrum band exceeds a threshold and this gives a good indication of how much that particular band is being used.

Whilst this approach is appropriate for bands that are dominated by one service type, it does not fully address the question of the level of problems being encountered in the LE bands, which are characterised by many, varied types of service. In these bands the issue becomes one of looking for problems and this is normally expressed using the word 'congestion'. The working assumption is usually that, as the LE bands become more heavily utilised, more problems will occur and users will experience loss of service quality. This project started out with this assumption, but decided that it is misleading.

3.2 HIGHER LAYER MONITORING

In Figure 11 there are two questions illustrated against the main protocol layers of interest to this study. This diagram summarises the common understanding at the start point of this project. For simplicity the protocol stack layers have been kept to the physical layer, the wireless link layer and then all higher layers. In a heterogeneous WiFi network the actual protocol stack is considerably more complex than this, but this model will suffice for explaining the concepts being considered here.

	Utilisation How busy is it?	**Congestion** How many problems are there?
Higher layers	Packet rates	Packet errors
Wireless link layer	Frame rate Throughput Number of nodes	Retry rate Frame loss rate AP queue length
Physical layer	Spectrum utilisation	

Figure 11 Utilisation measurement - initial model

The first question is that of utilisation. Measurements of utilisation try to estimate how busy a network is and can be performed at any layer in the protocol stack. As mentioned above, physical layer monitoring is the conventional approach. This has the advantage of producing results that are completely generic and not dependent on radio modulation, protocol or service type.

The second question in Figure 11 is that of congestion. Whilst the word 'congestion' is in common usage, we are more interested in measuring how many problems there are. In fact we are actually interested in measuring the user experience that can be expected in an environment.

Looking for network problems cannot be readily performed at the physical layer, as it is not possible to see frame and packet errors without demodulating the signals. It is far easier and more accurate to accept that the signals must be demodulated and then analyse the data frames/packets.

In general it is better to perform monitoring as low down in the protocol stack as possible. This minimises issues of data confidentiality and reduces the complexity and number of demodulators required.

We are led to the preferred model for measurement shown in Figure 12. This says that, if utilisation measurement is sufficient, then spectral utilisation measurement at the physical layer is the best approach. If, however, the user experience is to be understood, then the degradation also has to be measured and it is preferable to work within the wireless link layer.

After analysing the survey data we concluded that the frame and retry rates are the best parameters to use. These are parameters that were found to correlate best with the utilisation and degradation respectively.

	Utilisation How busy is it?	**Degradation** What service quality can the user expect?
Higher layers		
Wireless link layer	Frame rate	Retry rate
Physical layer	Spectrum utilisation	

Figure 12 Utilisation measurement - preferred model

In this new model we have deliberately removed the notion of measuring 'congestion' as the word is not a helpful one in this context. Congestion is just one of a number of factors that can lead to network degradation.

In order to clarify this concept of network degradation, section 4 discusses it in detail and presents a taxonomy of possible problem classes.

4 NETWORK DEGRADATION

This section discusses the many and varied types of problem that can occur when using WiFi to access the Internet, particularly in urban areas. A literature search has revealed a wide range of documented research into networking problems. There are also many undocumented problems and these have been investigated by consulting users and industry stakeholders. Taking all these into account, a taxonomy is given to help categorise these problems.

4.1 NETWORK STRUCTURE

It is helpful to have a model of network structure in mind when thinking about networking problems. A simple model is shown Figure 13 in which a client device (e.g. laptop PC) connects to a server computer via an Access Point (AP) operating in infrastructure mode. Many other configurations are possible, but this model is sufficient for the purpose of the current discussion.

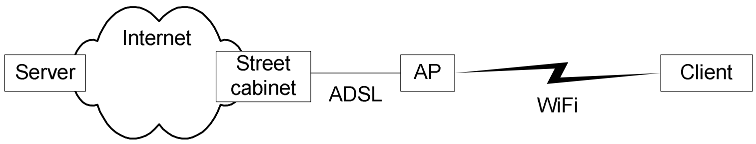

Figure 13 Simple model of Internet access via WiFi

Connecting to the Internet via WiFi is a complex business and one which requires the successful configuration and ongoing operation of a number of communications protocols.

As a simple example, Figure 14 shows the protocols that are involved in displaying a GIF image in a browser. Each protocol in layers 2 to 5 must work correctly and efficiently for the user to have a good experience. Also, the physical layer[1] must provide reliable connectivity via the shared 2.4 GHz ISM band.

[1] The IETF (1989) specifies layers 2 to 5 for the Internet. Layer 1 is added here to distinguish the shared physical medium from the link layer protocols.

	Server	Internet Routers		Street Cabinet		AP		Client
5 Application layer	IMAGE-GIF							IMAGE-GIF
	HTTP							HTTP
4 Transport layer	TCP							TCP
3 Internet layer	IP	IP		IP		IP		IP
2 Link layer	Wired	Wired	Wired	Wired	ADSL	ADSL	LLC	LLC
							WLAN	WLAN
1 Physical layer	Wired		Wired		Phone line		2.4 GHz ISM band	

Figure 14 Example of protocol stacks

Note that Figure 14 rather simplistically refers to the 'wired' network, which is sufficient here as the wired physical layer tends to be considerably more reliable than the wireless physical layer. Note also that Asymmetric Digital Subscriber Line (ADSL) is assumed as the access method.

When a person connects to the Internet via WiFi network and uses, for example, a web browser, each command sent from the client to the server and every page on the server to be displayed on the client must be transferred through a series of protocols similar to those depicted in Figure 14.

Errors can occur at any point in this chain, not just in the WiFi protocols (bottom right of diagram). The user rarely gets a clear indication of where an error has occurred and most errors are never reported to the user directly.

4.2 NETWORK DEGRADATION TAXONOMY

Figure 15 shows the various classes of degradation that can occur arranged as a taxonomy tree[2]. Problems that affect the user can originate in the physical layer, the link layer or the higher layers.

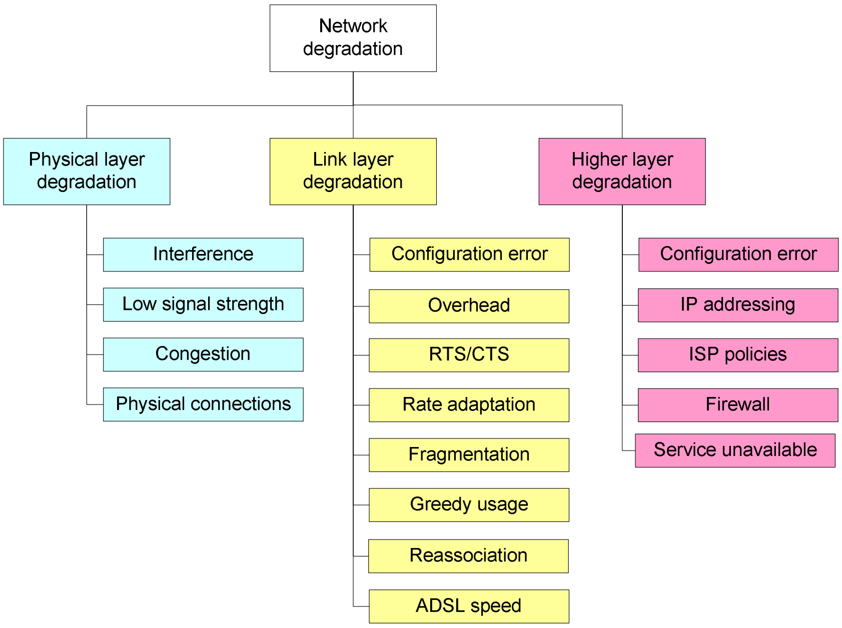

Figure 15 Taxonomy of degradation classes

The following paragraphs describe these different classes of problems and give some examples. The objective is to illustrate the main concepts rather than present an exhaustive list of problems.

It will be seen from these lists that the sources of problems are many, varied and often complex. Few people will be aware of all of these and, because of the information available to them, may tend to attribute their own communications problems to high levels of load offered to the network. They therefore tend to use the term 'congestion' to describe their experiences irrespective of the cause of their specific problems.

[2] Note that the colours used in Figure 15 are the same as those used in Figure 11, Figure 12 and Figure 14.

4.2.1 Physical layer degradation

Degradation Class	Description
Interference	The IEEE 802.11 family of WiFi protocols (IEEE, 2007) has been designed to work together, but the 2.4 GHz ISM band is shared with many other types of service. Many devices will interfere with the WiFi signals causing complete or partial blocking of one or more channels (Cisco, SA). Moving to a different channel will sometimes alleviate the problem (De Maesschalck, 2006). e.g. interference from (Cisco, SA), (Cox, 2007) (Mitchell, SA): • bonded channel interference* • microwave oven leakage (section 7.5.3); • audio video senders (sections 7.5.3 & 7.6.1); • baby monitors; • cordless phones; • garage door openers; • home automation devices; • fluorescent lights; • game consoles and wireless game controllers; • other communications (Bluetooth, Zigbee, etc.).
	* A potentially significant future source of interference is the incompatibility between IEEE 802.11b/g networks and the implementations of IEEE 802.11n that use two channels bonded together to increase the data rate. Higgins (2007) states that the bonding of two channels in IEEE 802.11n is likely to cause interference with existing IEEE 802.11b/g networks. Intel (2008) proposes that channel bonding is used only in the 5 GHz band to avoid problems at 2.4 GHz.

<table>
<tr><th>Degradation Class</th><th>Description</th></tr>
<tr><td>Low signal strength</td><td>Users try to operate this short-range technology at ranges that exceed its capability, leading to connectivity problems, e.g.:
Loss of connection at long range;
Frequent IP address renewals due to reconnection.
Many devices are now on sale that extend the usable range showing that there is a demand for long-range links (Goodwin, 2007) (FreeNet, 2008) (Purdy, 2007).

There is also an impact due to the fading of the channel. Channel effects such as multipath mean that users will experience signal strengths that have drastic and random fluctuations (Pätzold, 2002, p.5) (Awoniyi and Tobagi, 2004).

Kemerlis et al (2006) showed that, when TCP is being used, clients with lower signal strength get less bandwidth than those with high signal strength.

Users are advised to move client devices closer to the AP (De Maesschalck, 2006) (Nintendo, SA).</td></tr>
<tr><td>Congestion</td><td>When a lot of traffic exists on one or more channels, then the congestion mechanisms built into the network protocols will limit the throughput achieved.

Whilst there is 83 MHz of spectrum available in the 2.4 GHz ISM band, WiFi only uses part of this and it is mainly channels 1, 6 and 11 that are used. Typically these channels are not utilised equally so one channel will normally be used more than the others (but note that our results do not suggest that this is a major problem, see section 8.4). Users are frequently advised to select a different channel if they are experiencing any kind of WiFi problem (Mitchell, SA) (Nintendo, SA).</td></tr>
<tr><td>Physical connections</td><td>The physical connections between networking components can sometimes cause problems (Phifer, 2007) and, as these problems can be intermittent, may be interpreted by users as wireless networking faults.

e.g.
WiFi router connection to ADSL line damaged;
Missing, faulty or poor quality microfilter at phone line socket (adslnation, 2009);
Damaged Ethernet cables.</td></tr>
</table>

4.2.2 Link layer degradation

Degradation Class	Description
Configuration error	Networking devices have many different configuration parameters. Incorrect setting of any individual parameter can lead to lack of access or loss of service quality at the link layer. e.g. • Incorrect security code (e.g. WEP key) (Mitchell, SA), (Phifer, 2007); • Different security protocols specified on AP and client (Phifer, 2007); • MAC filtering prevents access of client to AP (De Maesschalck, 2006) (Mitchell, SA) (Nintendo, SA) (Phifer, 2007); • Newer IEEE 802.11 mode not supported by older clients (Nintendo, SA) (Phifer, 2007); • Incompatibility of client devices with non-standard modes such as 'Turbo Boost' and 'g Nitro' modes (Ven, 2009), (Nintendo, SA), (Phifer, 2007); • Beacon interval incorrectly set (Linksys, 2009) • Transmit power set too low for required range. • Preamble type in AP doesn't match that in client (Bruce, 2002, p. 178)
Overhead	Users may expect to obtain data rates similar to those advertised by manufacturers and close to those suggested by the standards. The actual data rate may well be considerably lower than expected, which can be perceived as a problem. A substantial amount of the available data rate is used by the WiFi overheads (Mitchell, SA). This is confirmed by our measurements which suggest that normally 10% or less of the throughput is carrying user data (Section 7.4).

Degradation Class	Description
RTS/CTS	The Request To Send/Clear To Send (RTS/CTS) link layer protocol is an optional protocol that controls access to the shared physical layer by preventing a node from transmitting a data frame until an RTS/CTS handshake has been completed (Geier, 2002). It is intended for use when some network nodes are hidden from each other and requires an increase in overhead for the handshaking to take place, so can increase the overall load offered to the network. RTS/CTS must be configured correctly if it to be used successfully. This is achieved by setting a threshold on the size of frame that requires a handshake and most vendors recommend a threshold of about 500 bytes (Geier, 2002), (Linksys, 2009). Jun *et* al (2003) calculated the reduction in bandwidth efficiency due to using RTS/CTS. Jardosh *et al* (2005a) and Kemerlis *et al* (2006) found that the use of RTS/CTS by a few nodes in a congested IEEE 802.11b network can prevent those nodes from gaining fair access to the channel. Hence it must be concluded that RTS/CTS can solve the hidden node problem but probably tends to reduce throughput (Note, however, that Franceschinis *et al* (2005) disagreed with other researchers and said that it improves thoughput.)
Rate adaptation	The WiFi family of protocols is designed so that the data rate between any pair of nodes is continually changed to maintain an acceptable error rate. The criterion for this decision is based on the signal-to-noise ratio, so that the data rate reduces with decreasing signal strength. Heusse *et* al (2003) and Jardosh *et al* (2005a, 2005b) found that the logic of the rate adaptation mechanism in IEEE 802.11b is counter to that required in congested environments. It tends to decrease the data rate when congestion occurs, causing the frames to increase in length and hence exacerbating the congestion problem by reducing throughput.
Fragmentation	Fragmentation is a mechanism for passing long frames over unreliable communications links by splitting them into multiple, smaller frames. Changing the fragmentation threshold affects network performance and therefore the user experience. The choice of threshold is a balance between error rate and throughput (Phifer, 2004), (Linksys, 2009).

Degradation Class	Description
Greedy usage	Some devices transmit disproportionately high rates of link layer frames. These will increase the probability of frame collision leading to the need for frame retransmission. It is possible for a client close to an AP to deny usage to other clients by flooding the network with frames. In the extreme this becomes a Denial of Service (DoS) attack. e.g. • Rapid transmission of null data retry frames (see section 7.6.2). These frames appear to be transmitted by devices that are trying to go into their power save modes; • Short beacon intervals (normally 0.1024s, but smaller values, e.g. 0.01024s, are frequently observed, see section 7.6.3).
Reassociation	Whenever a client disconnects from an AP and reassociates there is an impact on the application layer (Raghavendra *et al*, 2007). Such events can occur very frequently in busy networks generating additional management overhead traffic and exacerbating any existing congestion problems.
ADSL speed	The speed of ADSL access is an area of heated debate currently in the UK and Ofcom (2008a) has produced its own report into the subject. There are a number of web sites offering speed tests (e.g. http://www.broadbandspeedchecker.co.uk/, http://www.broadband-expert.co.uk/broadband/speedtest/, http://www.thinkbroadband.com/speedtest.html) The data rate available over an ADSL depends on many factors, with contention being a significant cause of user perceived diurnal performance variation. The distance from the exchange has a very large effect on the speeds that users can expect to obtain (Ofcom, 2008a).

4.2.3 Higher layer degradation

Degradation Class	Description
Configuration error	Networking devices have many different configuration parameters. Incorrect setting of any individual parameter can lead to loss of service quality in the higher layers either within the wired Internet or in the wireless access network.
IP addressing	A common problem in WiFi networks is the use of DHCP to dynamically assign an IP address to a client device (De Maesschalck, 2006), (Phifer, 2007). Users with problems of this kind are normally advised to use static IP addressing. e.g. • Roaming aggressiveness set to a value that demands IP address changes that interrupt the application layer services. • DHCP lease time set to 'forever' is not recognised as a valid setting by client, which then constantly requests new IP addresses.
ISP policies	Debate continues over the extent to which ISPs should impose access restrictions to different services. Yoo (2006) argues for basing network management on congestion, particularly on bandwidth-intensive applications. Limiting of access speed by an ISP will be seen by a user as a network problem.
Firewall	A firewall comprises a packet filter and an application gateway (Tannenbaum, 1996, p. 411), both of which must be configured correctly in order to achieve the required access for users and maintain the required levels of security. Incorrect settings of a firewall can prevent application software performing the service required by the user (Phifer, 2007).
Service unavailable	Most Internet applications require a service to be provided by a server somewhere that is not under the user's control. Web servers, email servers, instant messaging servers, etc. all provide a service that user's expect to be available on demand. Problems at the server end of the connection will be seen by the user as a lack of service or reduced quality of service.

4.3 SUMMARY

This section has explored network degradation from different perspectives and shown that there are a wide variety of problems that can affect users trying to access the Internet via WiFi.

A taxonomy of network problems has been presented. This shows how varied the potential problems are, which is something that the majority of users will not be aware of. Many problems are related to configuration errors, which can cause loss of service quality with or without loss of connection.

Users naturally tend to attribute their own problems to congestion, but it is highly likely that most of these problems are actually due to other causes such as problems in the wired Internet or interference from other equipment operating in this band.

5 FIELD SURVEY APPROACH AND LOGISTICS

This section describes the general approach to carrying out the field survey and also discusses the logistics of carrying out such monitoring. The results of the survey are discussed in Section 7, Section 8 and Appendix C.

A number of approaches to measuring spectrum utilisation exist, broadly divided into active and passive monitoring methods. Passive monitoring has the advantage of not affecting the networks, unlike active monitoring, which stresses the networks by passing known data across them and looking for errors. Passive monitoring is generally easier to perform but produces less detailed information than active monitoring.

If we constrain ourselves to passive wireless-side monitoring, then this places some constraints on the information that can be obtained. However, in this way it is possible to determine a wide range of statistics about the network performance (Mahanti *et al*, 2007).

The approach for this study has been to carry a number of handheld devices (section 5.1) for several weeks in different places around the country. Monitoring concentrated on town and city centres as these were expected to be the places where problems are occurring (section 5.2). The detailed results are given in Appendix C.

Section 5.3 discusses our findings with respect to the logistics of performing passive monitoring in this way.

5.1 RECEIVER

The Nokia N810 Internet Tablet (Figure 16) was used for this task. These units have the advantage of having integrated Global Positioning System (GPS) receivers.

Figure 16 Nokia N810 Internet Tablet

The key specifications of the N810 are as follows:

Physical	Weight: 7.97 oz Length: 2.83” Width: 5.04” Thickness: 0.55”
Processor	TI OMAP 2420, 400 MHz
Operating System	Maemo IT 2008
Memory	DDR RAM 128 MB Flash 256 MB
Storage	2 GB internal memory Supports up to 8GB miniSD or microSD cards
Communications	Bluetooth v2.0 WiFi IEEE 802.11b/g USB 2.0
Other features	Built-in GPS QWERTY keyboard VGA web camera
Battery life	2½ hours (6 hours with external battery pack)
Monitoring software	Kismet Tshark Database

5.2 LOCATIONS

Monitoring was carried out in the following towns and cities. The populations of the various sites are taken from the Office for National Statistics (2001) and Aberdeen City Council (2008). Appendix C gives a series of maps showing the coverage at each location together with the key statistics for each 1 km grid square.

Town/ City	Population
London	**8,278,251**
Bournemouth (Dorset)	383,713
Derby (Derbyshire)	236,738
Aberdeen (Aberdeenshire)	209,260
Milton Keynes (Buckinghamshire)	184,506
Newport (Gwent)	139,238
Peterborough (Cambridgeshire)	136,292
Cambridge (Cambridgeshire)	**131,465**
Cheltenham (Gloucestershire)	98,875
Bedford (Bedfordshire)	82,488
Kettering (Northamptonshire)	57,803
Welwyn Garden City (Hertfordshire)	43,512
Hitchin (Hertfordshire	33,352
St. Neots (Cambridgeshire)	**27,372**
Huntingdon (Cambridgeshire)	20,600

Monitoring concentrated on the three areas highlighted, namely Central London, Cambridge and St. Neots. This gave a good view of three very different types of conurbation.

5.3 LOGISTICS

Different modes of transport were used for the monitoring including walking, cycling, driving and using the train. After analysing the data it was clear that at least 15 minutes of data is needed if the statistics are to be meaningful and an hour is desirable. This fact affects how such surveys should be carried out.

- **Walking**. The need to get 'into' the urban environment meant that walking was, by far, the most effective way of carrying out the monitoring. Even at walking speeds it is possible to walk through part of a 1 km grid square in less than 15 minutes, as seen in the Central London data, where the Tower Bridge area was walked through but not for long enough (see the map in Appendix C).

- **Cycling**. Cycling was effective but a little restrictive in pedestrianised areas of city centres. It is very good for housing estates, where there are typically a lot of small roads and cul-de-sacs in a small area. It is hard to drive in and out of these areas.

- **Driving**. This was less physically demanding than walking or cycling and had the advantage of a readily available power source (via the 12v sockets). Driving did not, however, allow the spectrum to be monitored very close to buildings in all urban areas.

  When driving there was tendency to drive along a lot of roads and cover a large area without dwelling within an area for long enough to get sufficient statistics. This can be seen clearly in the maps in Appendix C (e.g. Peterborough) where potentially useful data was discarded as there was less than 15 minutes of recording per 1 km grid square.

- **Train**. None of the data from train journeys contributed significantly to any of the results. Trains are not appropriate for this kind of monitoring because train speeds are normally too high and they do not travel within range of many networks.

Two main approaches to monitoring were tried as part of this research:

- *Ad hoc* monitoring:

  In this style of monitoring people were given receivers to carry as they go about their daily business. There was no control on where the people went after initially selecting them based on where they live and work. With a large number of relatively inexpensive receivers, a large survey could be carried out very cost-effectively in this way.

- Targeted monitoring:

  Here people were dedicated to surveying a particular location. This is a more effective approach when the number of receivers is small as all the data obtained is relevant and useful. Targeted monitoring is relatively expensive compared to *ad hoc* monitoring but produces relevant results more quickly.

The preferred approach for future surveys would be to use a mix of targeted and *ad hoc* monitoring with the emphasis on walking or cycling in urban and busy areas. When carrying out targeted monitoring it is important that sufficient time is spent in each grid square. For the purposes of future survey planning, an average of five grid squares per day is a reasonable target.

5.3.1 Resolution

The limit on GPS performance for these receivers is dictated by the behaviour when used indoors. Providing GPS lock is achieved before entering a building they will continue to produce position estimates once indoors.

Figure 17 shows the typical behaviour indoors from which it can be seen that the position estimates tend to jump around within a radius of about 50m of the actual location (shown as a red dot).

Figure 17 Indoor GPS behaviour

We found that a resolution of 1 km was appropriate for the purposes of this exercise. A smaller resolution would make it difficult to gather sufficient data and could lead to difficulties with the privacy issues that may be associated with publishing data identifying specific sites. A larger resolution would not allow the locations of trouble spots to be identified. We used the standard British National Grid (BNG) naming convention for 1 km grid squares, which will facilitate comparison with any future survey results.

The accuracy of the built-in GPS receivers was found to be more than sufficient for monitoring at a 1 km resolution. Looking at the graphs of routes in Appendix C it is clear that the accuracy is good enough to identify which roads were being followed. There are some large errors associated with the GPS signals being lost (see the example on the Derby map) but the effect on the results is only significant when travelling by train or using the London Underground. In most cases the accuracy is in the order of a few tens of metres, even in very built-up areas.

5.4 SUMMARY

This section described the approach to performing the field survey, which used small, handheld devices with built-in WiFi and GPS receivers. These could be carried around easily in densely populated areas, including into buildings.

It was shown that the GPS accuracy of such devices is appropriate for this kind of survey and we have found that a 1 km grid square is a suitable choice of resolution for presenting the results of a survey.

Walking is the best way of performing such a survey as each grid square must be monitored for long enough to obtain good statistics. With shorter monitoring periods the variability in the data is too great and it is hard to see the trends across an area. We estimate that the minimum useful monitoring period is 15 minutes in each grid square and have discarded our data when the amount of data collected was substantially below 15 minutes.

6 LABORATORY EXPERIMENTS

This section describes the approach to carrying out laboratory experiments. The results are discussed in Section 7 and Section 8.

The field survey described in section 5 produced a large quantity of data, but, on its own, this was not sufficient to understand what is happening to the networks. A series of laboratory experiments was therefore performed with the aim of trying to establish what level of service a user might expect to experience in different network states.

6.1 EXPERIMENTAL NETWORK

An experimental network was assembled from a variety of different types of device. The aim was to observe the performance of a heterogeneous network running under different offered loads.

The network comprised various combinations of the following devices:

Device	Quantity
ADSL modem/ WiFi router	4
Laptop PC (Windows XP)	4
Laptop PC (Red Hat Linux)	1
Desktop PC (Windows XP)	1
Desktop PC (Debian Linux)	1
Nokia N810 Internet Tablet	5
USB WiFi ‘dongle’	6
PCMCIA WiFi card	1
Sony Playstation 3	1

The network activity was monitored using an additional Nokia N810 and all testing was carried out on WiFi channel 6.

As well as running the experimental network normally, the effects of interference were of interest. The following devices were therefore added to the environment when interference testing was required:

- Panasonic NN-435W microwave oven;
- Technika TT-07 AV sender.

6.2 EXPERIMENTAL METHOD

With such a complex network it proved very difficult to obtain repeatable results. The following method, based on the BBC live news channel, was eventually developed. This method produced results that could be repeated with a good degree of confidence and the results correlated well with the user experience obtained in different configurations.

1. Configure all network devices and set up the background traffic required. Simulate background traffic using flood pings and large file transfers;
2. Start Internet Explorer on the test PC (one of the XP laptops connected to Router 1) and clear its temporary internet files;
3. Start the BBC live news channel (www.bbc.co.uk/news) which has been found to be a reliable source of video traffic;
4. Wait until the video stream is running smoothly;
5. Start the monitoring software on the Nokia N810;
6. If required, start the interference device;
7. Start two stopwatches. It was convenient to use the freely available PC Chrono software for this;
8. Leave one stop watch to run for three minutes. Call this time S1;
9. Pause the other stopwatch every time the video is not running smoothly;
10. At the end of three minutes note the elapsed time on the second stopwatch. Call this time S2. The ratio S2/S1 is used later for estimating user experience;
11. Stop the monitoring software;
12. Note down the Mean Opinion Score (see section 8.1);
13. Transfer the recorded data from the Nokia N810 to the PC and document the results.

Typically about three of these tests could be run in an hour. Just over 300 recordings were made in different combinations of equipment and offered loads.

The test method was developed on the assumption that congestion due to heavy load would be the major issue. It transpired that this was less important than the effects of interference. Fortunately, the same test method could be applied to the interference cases as well as the busy network cases.

6.3 SUMMARY

This section described the experimental network and method used. These experiments used a heterogeneous network in various loading conditions and were designed to gather the information necessary to link the user experience with the technical measurements that can be made in the field.

7 ANALYSIS OF RESULTS

This section looks at the results from the field survey and laboratory experiments with a view to understanding the measurements made. We do not address the user experience quantitatively here, but concentrate on interpreting the technical measurements only. Extending the analysis to the effects on users is the subject of Section 8.

It will be shown that, whilst network performance tends to degrade in more populated areas, the laboratory experiments do not support the idea that congestion is the main cause of network degradation. This leads to the conclusion that network degradation is complex and involves many factors, as explained in section 4. Interference, in particular, is believed to account for the majority of physical and link layer problems encountered by users.

7.1 NETWORK AND POPULATION DENSITIES

It is relatively easy for statistics to be gathered on the network density and many people contribute to the Wireless Geographic Logging Engine (WiGLE) database at www.wigle.net. The database has been running for seven years and contains the locations of millions of APs worldwide. There is, therefore, a wealth of evidence to show that WiFi networks are concentrated in the urban areas. The first confidence check of our data is therefore to relate population density to network density.

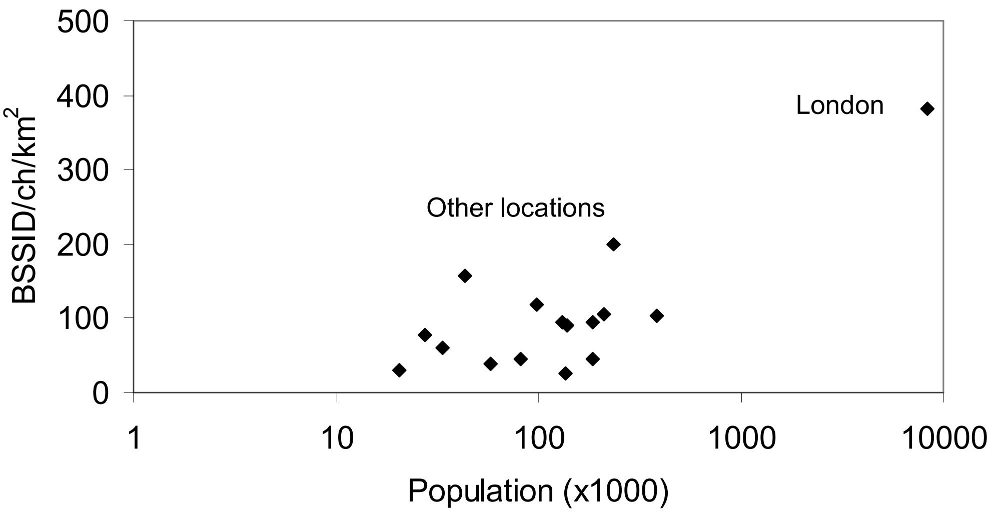

Figure 18 Network density versus population size

Comparing the network density from our field survey data with the population density shows that there are considerably more networks per square kilometre in London than the other locations.

- Figure 18 is a scatter plot of the median network density against the population for each location. The network density is expressed in terms of the number of BSSIDs per channel per square kilometre. The population axis is logarithmic in order to accommodate London's population, which is so much higher than the other locations.

- London clearly stands out in Figure 18 as having not only a much higher overall population than the other locations, but also a much higher network density. Note that this data is dominated by measurements from the centre of London rather than the suburbs.

The publicly available census information does not give population density statistics, so it is not possible to compare population density with network density at a smaller resolution.

The implication of Figure 18 is that there is direct link between network density and user experience. This is the concept of 'congestion', which asserts that a high network density implies a drop in network performance. This concept does not, however, fit well with our results. It is better to look at the traffic travelling through the networks rather than just the number of networks.

7.2 MEASURING UTILISATION VIA FRAME RATES

Many different parameters could potentially be used for describing how busy a particular location is. If one looks only at the link layer, then one can count the number of frames per second, the time intervals between frames, the number of bits in each frame and the duration of each frame. Each of these parameters is a candidate for measuring the utilisation of the band. It is also possible to consider each type of frame separately (e.g. beacon frames, user data frames).

Analysis of the data gathered during the survey and laboratory experiments suggests that the number of frames per second (i.e. frame rate) is the best parameter to measure. The frame rate produces the best correlations with other data, in particular the user experience metric, which will be described in section 8.

The frame rate is a very bursty parameter, varying rapidly over time and it is hard to define an instantaneous value for it. Appendix B looks at the choice of averaging period and concludes that an averaging period of five seconds is appropriate if one wants to estimate an instantaneous frame rate. We have looked at the statistics of the instantaneous frame rate for this report and concentrated on using the mean and maximum to describe an environment.

IEEE (2007) defines the IEEE 802.11 link layer in detail, but for our purposes it is sufficient to classify frames in a relatively simplistic way. The table below lists all the frame types defined by IEEE (2007). It will be seen that these fall into three major groups: Management, Control and Data frames and that subtypes exist within these groups.

Management frames	Control frames	Data frames
Association request Association response Reassociation request Reassociation response Probe request Probe response Beacon Announcement Traffic Indication Map Disassociation Authentication Action	Block acknowledgement request Block acknowledgement Power-save poll Request to send Clear to send Acknowledgement Contention-free period end Contention-free period end acknowledgement	User data Data + Contention-free acknowledgement Data + Contention-free poll Null data Null data + Contention-free acknowledgement Null data + Contention-free poll QoS data QoS data + Contention-free acknowledgement QoS data + Contention-free poll Null QoS data Null QoS data + Contention-free acknowledgement Null QoS data + Contention-free poll

IEEE 802.11 Frame types

With all these frame types it is hard to see the relevant information, so it is convenient to divide the frame types in a slightly different way. For the purposes of surveying the environment we have defined the following categories of frame type:

- Total frames (all frame types);
- Beacon frames;
- User data frames;
- Other management & control (all frame types excluding beacons and user data).

The total frame rate gives an indication of how busy the environment is. In section 7.3 the total frame rate is compared with the network density.

The beacon frame rate relates to the number of access points, although not all access points are configured to transmit at the same beacon frame rate. This parameter generally indicates the potential for utilisation in an area. A high beacon frame rate suggests a high network density.

User data rate indicates how much ('revenue-earning') traffic is being carried by the wireless networks.

The 'other management & control' category indicates how many frames are being used to keep the networks functioning correctly (apart from the beacon frames).

There are therefore four main categories of frame type used in this report. By comparing the relative proportions of these it is possible to gain some insight into how efficiently the networks are performing. In a perfect world, the user frame rate would equal the total frame rate, but practical considerations mean that the link layer has to carry non-user data frames so that the networks can operate in the desired way.

7.3 MEAN FRAME RATE AND NETWORK DENSITY

The network density and frame rates have been introduced in Sections 7.1 and 7.2. Both of these are potential indicators of how busy an area is, so the question arises as to which is better? It was asserted in section 7.2 that the total frame rate is our preferred parameter, but a brief comparison of the two is presented here for completeness.

Figure 19 plots the mean total frame rate against the network density for each of the 1 km grid squares surveyed in the monitoring activity (Appendix C). The correlation coefficient between these two variables is found to be about 0.73. This is not a particularly good correlation and is explained by the fact that not all APs are loaded equally. In a heterogeneous, urban environment, one expects to see considerable variation in the usage of WiFi networks and this is confirmed by Figure 19.

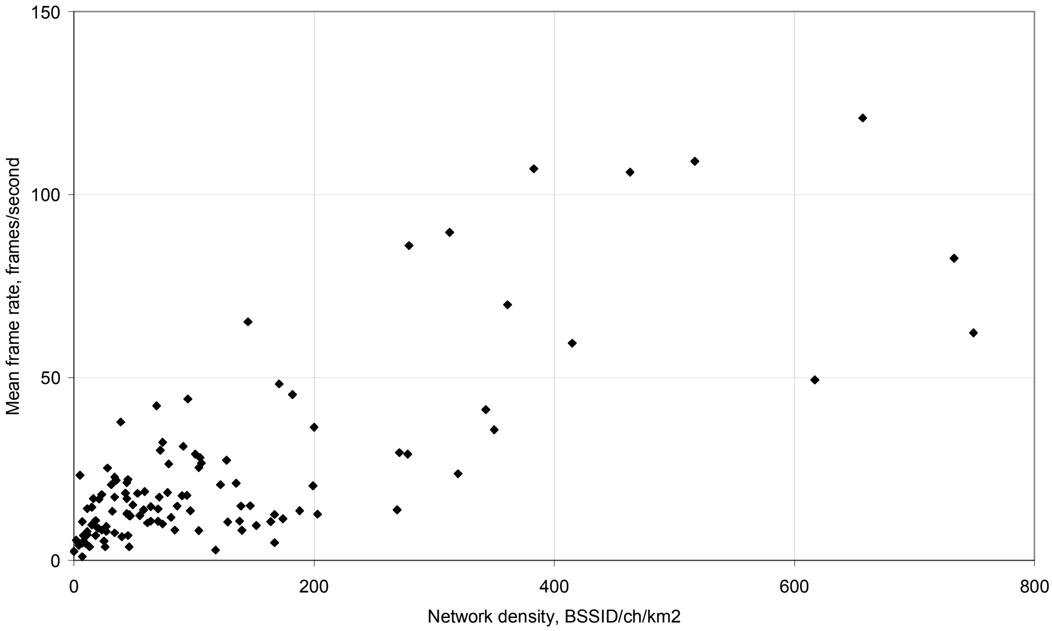

Figure 19 Mean total frame rate versus network density

We take the view that it is better to use the total frame rate than the network density as a measure of utilisation. It has the advantage that it can be readily compared with the retry ratio (Section 7.5) and it indicates the actual usage in any given area rather than the potential usage. The number of networks in an area is unlikely to change significantly throughout the day, but the actual utilisation will change.

On this assertion that the total frame rate is the better parameter to use for looking at network utilisation, we now consider the levels of utilisation seen at different sites.

7.4 UTILISATION

The density of WiFi traffic is considerably higher in London than the other locations surveyed, suggesting that the utilisation is correspondingly higher.

- This is clear from Figure 20, which shows the mean frame rate for each location. This was produced by calculating the minimum, median and maximum of the mean frame rate across each of the 1 km grid squares used in each location.
- In Figure 20 the locations are ordered by decreasing population size. The medians of the mean frame rates do tend to fall with decreasing population, although there is considerable local variation.
- London and Welwyn Garden City stand out as having particularly high mean frame rates.

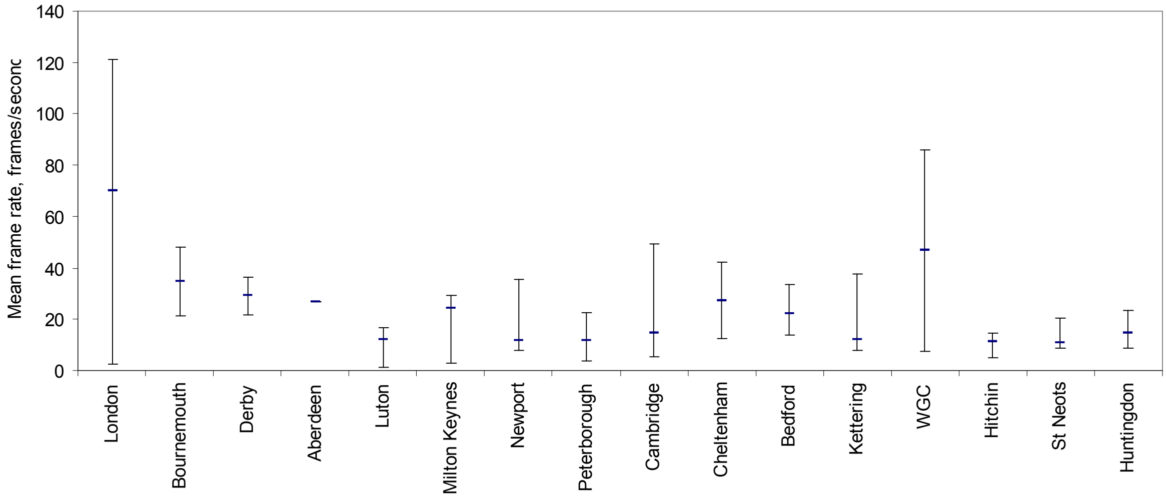

Figure 20 Frame rates in different locations

Figure 21 shows a breakdown of frames by type for some of the grid squares. These have been selected from the data presented in Appendix C by choosing the grid square with the highest mean frame rate. The figure illustrates the wide variety of utilisation characteristics in different areas.

- Generally the user data frames are less than 10% of the total frame rate. In London TQ 33 81, for example, the user data is almost exactly 10% of the total.
- Bournemouth SZ 07 93 is an exception to this, with 44% of the frames being user data frames, indicating well-utilised networks with very low levels of degradation.

- Welwyn Garden City (WGC) TL 23 13 has a relatively high mean beacon frame rate compared with London TQ 33 81. The WGC site has a network density of 279 BSSIDs/ch/km^2 which is much lower than the 657 BSSIDs/ch/km^2 seen at the London site. Also, the London site has APs with short beacon intervals (Section 7.6.3).

  We believe that some of the frames are being lost in London, because either interference or congestion[3] is occurring in London but not WGC. The results of our experiments (Section 7.5.3) suggest that interference is the more likely cause of this frame loss in London.

User data was a relatively low percentage of the traffic in most of the areas surveyed. The impression obtained is of a large background mass of messages needed to keep the networks alive. On average, a smaller quantity of user data is carried across the networks, but this is bursty so peaks of user traffic are sometimes seen.

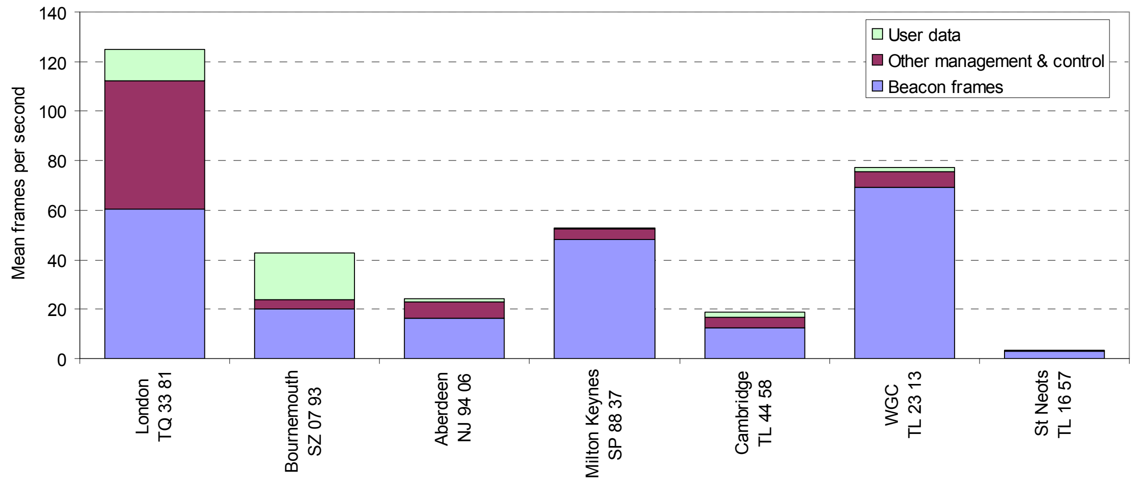

Figure 21 Proportions of frame types in selected locations

It would be useful to further categorise the frame types to a lower level of fidelity in an effort to identify mechanisms for degradation. However this analysis has not been attempted here as the effort involved was not considered justifiable in the context of this study.

As a final comment on utilisation, Kershaw (2004, p.18) states that, during one Denial Of Service (DOS) attack on a network the frame rate was 1700 frames/second. This places the above data into some degree of perspective.

[3] Jardosh *et al* (2005b) used the number of missing beacon frames as a measure of link layer reliability. We have tried to use this approach to measuring congestion, but found it unworkable in the urban environment where there is no control by the survey team over the network infrastructure.

7.5 RETRY RATIO

IEEE (2007) defines a number of frame types, as described in Section 7.2. It also allows management and data frames to be tagged as retries. A retry frame indicates that a frame has been lost at some point between transmitter and receiver. Looking at the number of retry frames is a useful measure of the level of problems being encountered by a wireless network.

The above analysis concentrated on the mean total frame rate, λ. In this section we use the retry rate, μ, which is normalised by dividing it by the frame rate. The normalised parameter is called the retry ratio, γ, i.e.:

$$\gamma_k = \mu_k / \lambda_k \qquad (1)$$

where k is the index of the 5 second averaging window.

The retry ratio defined above indicates the proportion of frames that are being lost. As such it is a useful measure of network problems without being overly specific about what those problems are. In a network that is working faultlessly the retry ratio will be zero. Conversely, a network that is experiencing a lot of problems will have a retry ratio approaching 100%. The retry ratio will vary over time, but, typically, real networks that are working well exhibit mean retry ratios of less than 10%.

In the next three sections the results of the survey and experiments are presented in terms of the mean and maximum values of the retry ratio. It will be shown that the retries seen in the survey data cannot be explained by simply increasing the WiFi traffic, but that interference does provide an explanation.

Sections 7.5.1, 7.5.2 and 7.5.3 present the results of the survey and laboratory experiments in a series of graphs. These graphs are plotted on the same axes so that the results can be compared with each other.

7.5.1 Survey data

Figure 22 shows the mean retry ratio plotted against the mean frame rate for all the data collected in the survey. It will be seen that the mean retry ratio tends to increase with increasing mean frame rate, but there is only a weak correlation between them. Even at very low mean frame rates the mean retry ratio can approach or even exceed 10%.

Figure 23 shows the same data but using the maximum retry ratio and maximum frame rate. The maximum retry ratio increases with increasing maximum frame rate, but there is not a good correlation between these parameters. Note that the maximum retry ratio does, at one site, reach 100% suggesting severe network degradation. That site is in St. Neots (TL18 62) which is not a very busy location.

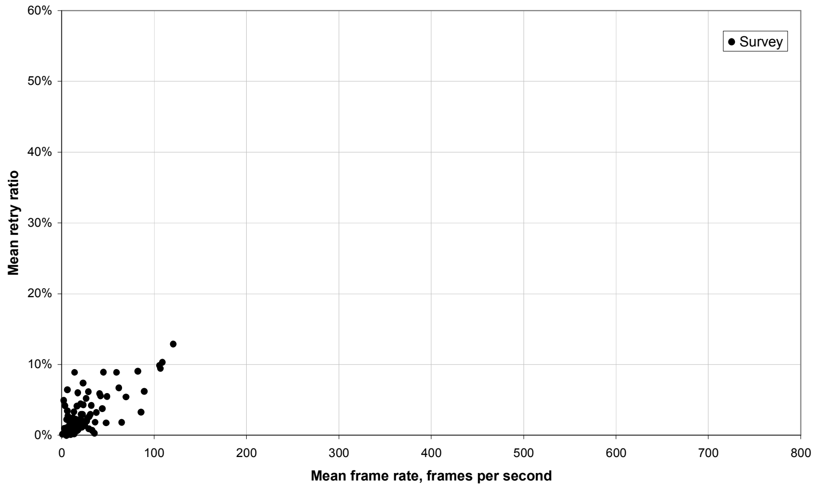

Figure 22 Mean retry ratio versus mean frame rate, all survey data

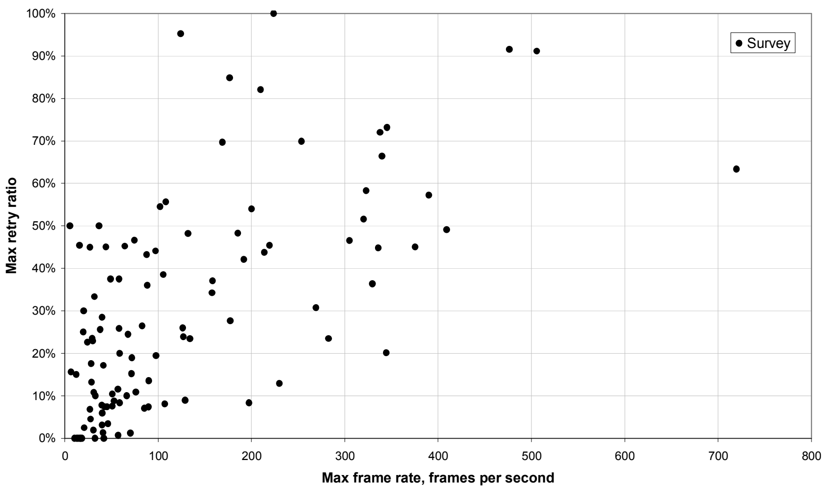

Figure 23 Maximum retry ratio versus maximum frame rate, all survey data

7.5.2 Experimental data

Figure 24 is similar to Figure 22 but shows the mean retry ratio versus the mean frame rate for the survey data and also the experimental data.

In the experimental data there is little increase in the mean retry ratio as the mean frame rate increases. In fact the mean retry ratio remains well below 10% in most cases. Hence, the experiments suggest that increasing the load on the networks does not, at the levels used here, significantly increase the number of retries needed.

Figure 25 shows the same data in terms of the maximum retry ratio and frame rate. Comparing Figure 25 with Figure 23 it is again clear that the levels of network degradation seen in the survey were not recreated in the laboratory.

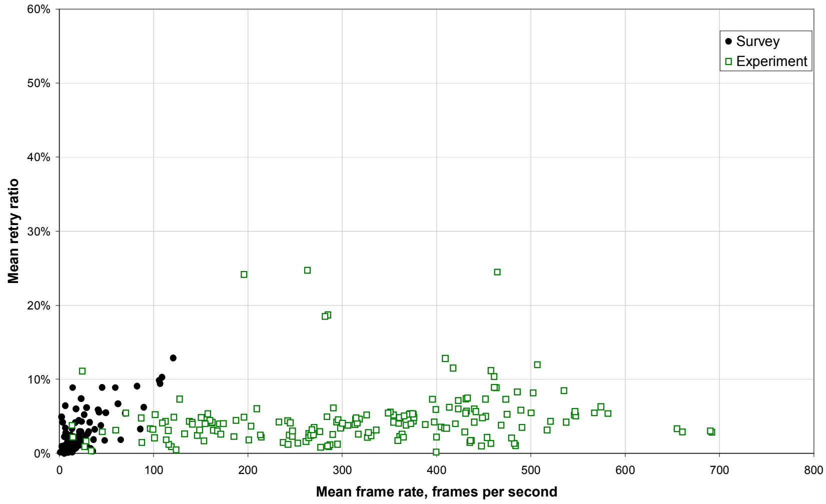

Figure 24 Mean retry ratio versus mean frame rate, survey and experiments

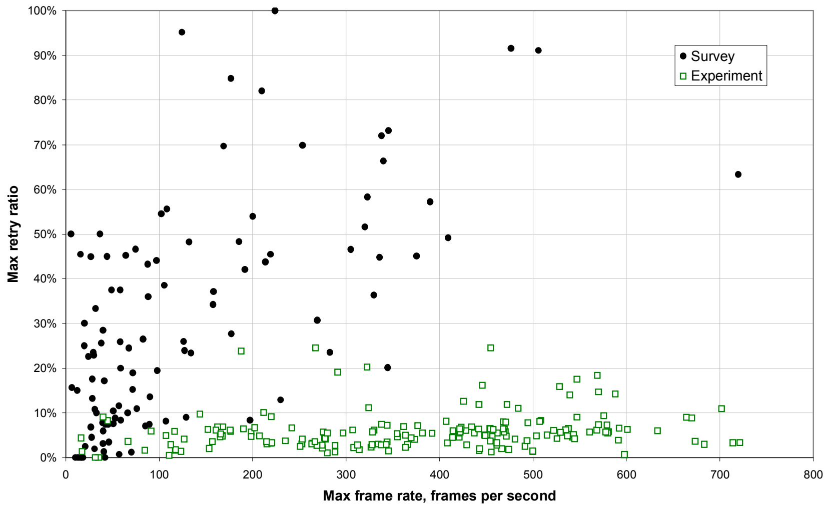

Figure 25 Maximum retry ratio versus maximum frame rate, survey and experiments

7.5.3 Interference data

Figure 26 shows the same information as Figure 24, but with points added for the experimental network with interference sources added (AV sender and microwave oven).

It is clear from Figure 26 that interference frequently causes a significant increase in the mean retry ratio, even at very low frame rates. From this we conclude that interference effects are the most likely to cause the elevated levels of retry ratio seen in some of the survey data.

Figure 27 is a repeat of Figure 25 with the interference points added. Again we conclude that it is only by adding interference that the levels of retry rate seen in the survey can be replicated.

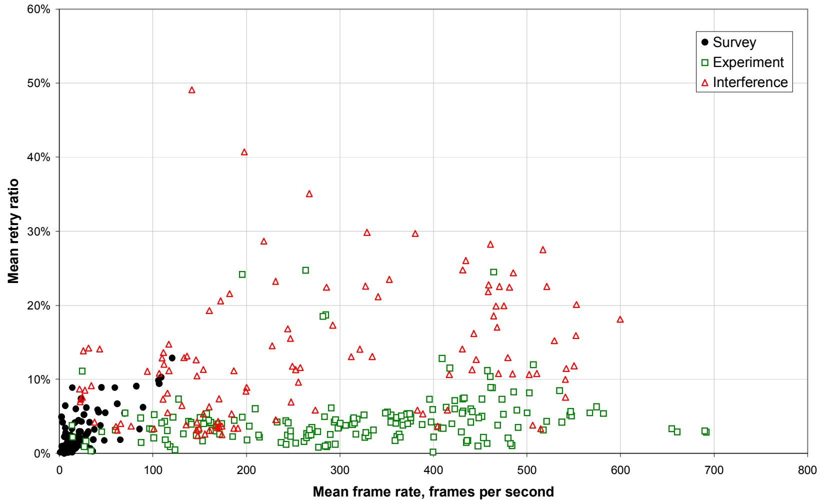

Figure 26 Mean retry ratio versus mean frame rate, interference added

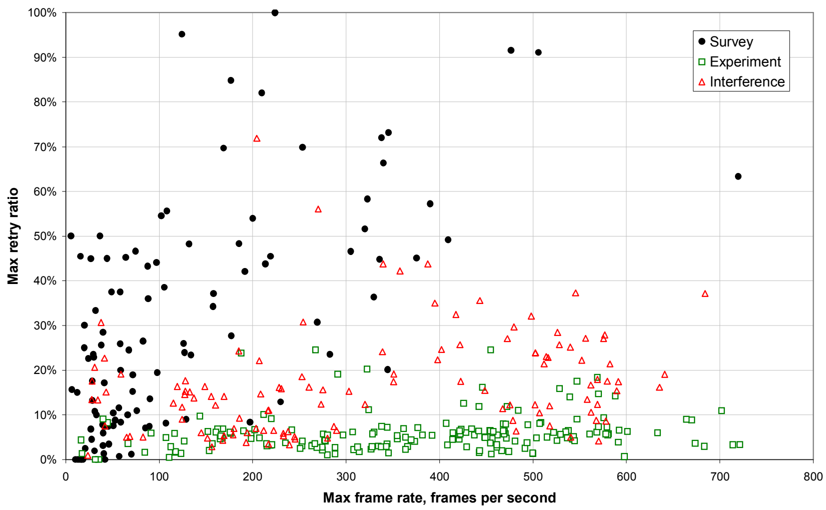

Figure 27 Maximum retry ratio versus maximum frame rate, interference added

The graphs presented above show that high levels of network traffic, as measured by the frame rate, cannot account for the severe degradation seen in the field, as measured by retry ratio. Rather, the network performance is seen to be extremely sensitive to interference, which leads us to conclude that interference is the dominant cause of degradation in WiFi networks.

7.6 NETWORK PROBLEMS

In section 7.5 it was concluded that WiFi network loading is not the main cause of network degradation and that interference is the most likely reason for seeing high levels of retry ratio in the survey data. In section 4.2, however, it was explained that a large number of factors can lead to network degradation and that interference is just one of many types of problem.

This section looks at some specific types of network problem that have been revealed during the laboratory experiments and during low-level analysis of the data.

7.6.1 Audio video senders

Several times in this report we have stated that audio video (AV) senders are a problem for WiFi services. This conclusion comes from a series of laboratory experiments investigating the impact of AV senders and from looking at how popular these devices are.

In section 7.5.3, the impact of microwave oven and AV sender interference was shown in terms of the increase in retry ratio (see Figure 26 and Figure 27).

In another experiment an audio video sender was operated whilst monitoring the packet loss on a WiFi link running on channel 6. The results are shown in Figure 28 for each of the four channels selectable on the AV sender.

When the AV sender was running on its channel 1 or 4, then there was very little impact on WiFi channel 6. However, when the AV sender was switched to its channel 2, then WiFi channel 6 was completely lost. If the AV sender was operated on its channel 3 setting, then about 50% packet loss was seen on WiFi channel 6.

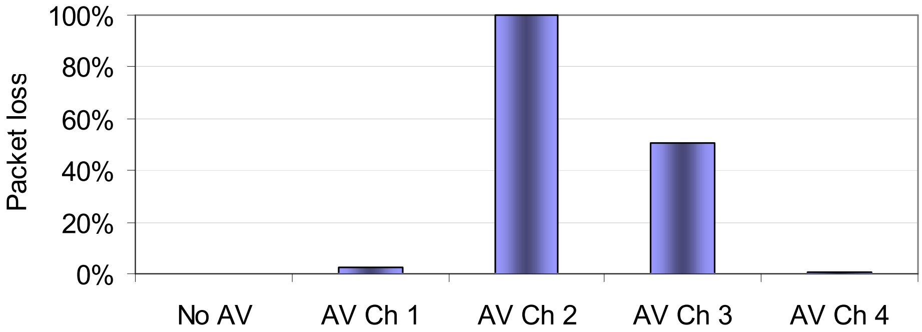

Figure 28 WiFi packet loss due to interference by AV sender

What is not shown in Figure 28 is what the user of the WiFi channel saw. When the AV sender was set to its channel 2, then the AP was clearly visible to the WiFi user and maximum signal strength was indicated. Despite being able to 'see' the AP, no connection could be established.

This effect is identical to that reported anecdotally by several people during the course of this study. It can easily account for the problems experienced by people who are blaming their neighbours for abusing WiFi by overuse.

Estimating the number of AV senders in use is not easy. X10 (SA) estimated that the European market at one million units in 2002. Li, *et al* (2008) quoted a J.P. Freeman estimate of the market for networked surveillance cameras in the USA as $600m in 2008.

There are many models to choose from, suggesting a healthy consumer market. There are three main markets: satellite TV, baby monitoring and home security.

Many AV senders are aimed at owners of satellite TV receivers who want to relay the TV around their houses. A short Internet search for 2.4 GHz devices reveals over twenty different models currently on sale.

Baby monitors are very similar devices, but packaged in a more attractive way for the task of baby monitoring than satellite TV viewing. A search for analogue 2.4 GHz devices produced eleven different models.

Another class of AV sender is that targeted at the home security market. The technology is the same as that used in the satellite TV and baby monitor, but the packaging is slightly different to allow for outdoor or covert use. Seven models of analogue 2.4 GHz devices were found in a short Internet search.

There are, therefore, over forty different models of AV sender currently on sale. All these, irrespective of the way they are marketed and packaged, use an 8 MHz bandwidth FM transmission to carry PAL TV and they will probably interfere with WiFi services on the same channel. They will also be interfered with by WiFi services. Normally this is not a major problem for the purchaser, as they can choose to put their domestic WiFi on a different channel to the AV sender. Their neighbours may, however, be adversely affected and not know what has caused the problem.

Some newer baby monitors claim to be WiFi-friendly. This suggests that the interference problem is being recognised and dealt with by some manufacturers. Six models of digital baby video monitor were found to be on sale in the UK.

As digital TV sets replace analogue ones, it is reasonable to expect this new generation of digitally modulated AV senders to replace the FM versions. Baby monitors seem to be leading the way in this regard, one reason being that people want to secure their baby monitors so that the signals cannot be viewed by anyone else, which is functionality that can be achieved using digital modulation. The digital AV senders tend to be more expensive than the analogue senders, but the additional features for the baby monitoring market mean that the additional retail price can be accommodated.

The home security market seems to be adopting IEEE 802.11g WiFi for the newer products. At the time of writing, only three devices were found that used this technology.

So, it is seen that there are a lot of products being sold to transmit video signals over the 2.4 GHz band. The majority of those currently on sale use analogue modulation and interfere with WiFi communications. We are, however, seeing a move towards digital modulation, which will be more 'WiFi friendly'. The baby monitoring market seems to be adopting an FSK modulation, whereas the security market is moving towards IEEE 802.11g.

The domestic TV AV sender market doesn't appear to be moving towards digital modulation yet. These devices are cheap and meet people's requirements, especially for distributing satellite TV signals from a central decoder. Interference to and from WiFi may be the stimulus for digital modulation to appear in this market.

7.6.2 Null data retry frames

Figure 29 shows an example of a recording made in London in which there is a ‘spike’ of points which are at high frame rate and high retry rate. In this graph, each point represents one second of the recording.

On closer investigation, it appears that all these frames are null data retry frames. Our interpretation of these events is that these frames are transmitted by a client device when it wants to enter power save mode (Putman, 2005). The WiFi protocols require a substantial amount of processing, so portable client devices will try to put their WiFi interfaces into power save mode in order to improve battery life.

A portable client will request that its AP stops forwarding frames and buffer them until the device is again ready to receive them. It appears that the null data retry frames are used for this purpose. If the AP does not acknowledge receipt of the null data retry frames then the device continues transmitting them; one assumes that the device will eventually stop transmitting.

In dense places like the centre of London, there can be a lot of devices exhibiting similar behaviour at the same time, hence there is a noticeable effect on the overall frame rate.

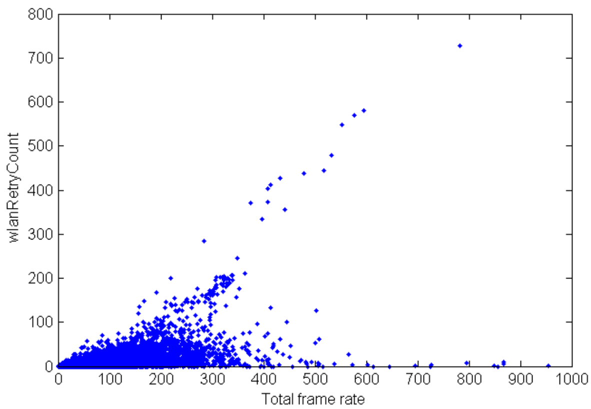

Figure 29 Client device trying to enter power save mode

The question then arises as to what proportion of the overall retry rate can be attributed to this effect. Figure 30 shows the percentage of retries that are null data frames and it will be seen that there is an increase in this percentage as the retry rate increases. In fact the proportion rises to nearly 80%, which suggests that this is a dominant effect.

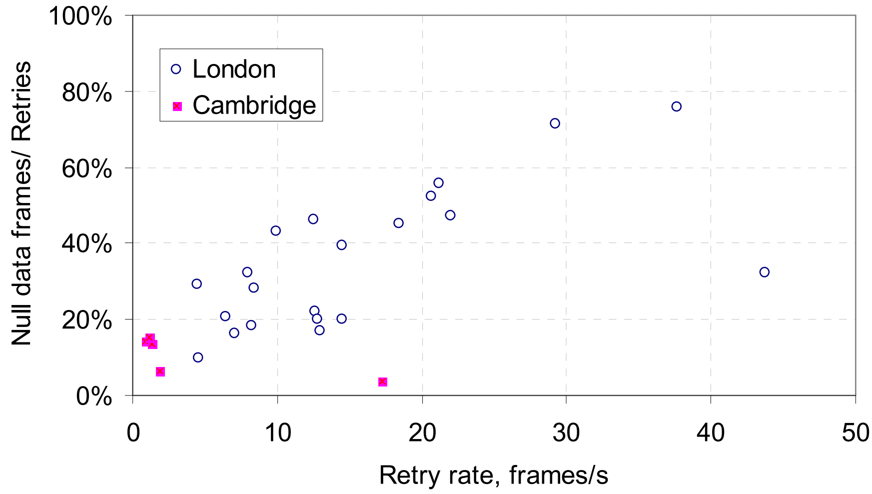

Figure 30 Percentage of retries that are null data frames

The implication of this investigation is that a significant proportion of the lost frames in the London environment are occurring during events when devices are losing contact with their access points and are trying to enter power save mode. A more parsimonious messaging protocol would therefore help to reduce the lost frames in the urban environment.

7.6.3 'Free Public WiFi virus'

This is a phenomenon that is regularly seen in the data collected in London and has become known as the 'Free Public WiFi virus' when reported on the Internet (Frary, 2008), (Kaleem, 2007), (Mac-net, SA), (Srivastwa, 2008).

What happens is that someone looks for available networks on a display something like Figure 31 with an SSID called 'Free Public WiFi'[4]. The person then tries to connect to it, but nothing happens. Such a network will be an ad hoc rather than infrastructure network and the user will be prompted whether to store it as a setting on their device. If the person confirms this prompt then their device will be 'infected' as it now starts advertising its presence as an ad hoc network.

Figure 31 Free Public WiFi virus display in Windows

Kaleem (2007) and Srivastwa (2008) attribute the problem to the Windows wireless auto configuration utility, but our tests suggest that the Maemo operating system behaves in a similar way, so it is likely that other operating systems will also propagate these fake SSIDs.

Having a few such networks in an area would not normally be a problem, but we have found that the beacon interval for these connections has been set to 0.01204 seconds rather than 0.1024 seconds. This is probably an installation error, but it does cause around a hundred beacon frames per second to be transmitted by the 'infected' device. Each device is therefore transmitting beacon frames at a rate equivalent to ten APs.

Walking around central London one can expect to encounter, on average, 26 of these devices per hour. On average they constitute about 7% of the total frame density, but this is very much an average and this peaks at about 28% in some cases. These devices are therefore contributing noticeably to the background utilisation, but are not, on their own, causing frames to be lost.

7.7 SUMMARY

The data collected in the field survey and our laboratory experiments has been analysed to gain an understanding of the main behaviours of heterogeneous WiFi networks in urban areas.

The total frame rate and retry frame rate have been found to be the best parameters to use for quantifying the level of utilisation and network degradation respectively. These parameters map well to the questions posed in Figure 12 (i.e. 'how busy is it?' and 'what service quality can the user expect?').

Three specific network problems have been discussed as examples of real-world effects that are commonly encountered. The most important of these is believed to be the impact of AV senders, wireless security cameras and baby monitors on WiFi services.

The results of the laboratory experiments show that WiFi networks are very sensitive to interference from other types of signal. Other effects, such as null data retries and WiFi viruses, are believed to contribute to network congestion, but their impact on the user experience is probably much lower than that of interference.

[4] Other common network names found in our data are 'Free Internet Access' and 'hpsetup'

8 USER EXPERIENCE

In section 7 the results of the survey and laboratory experiments were presented in terms of the frame rates and retry rates observed. In this section these technical measurements are related to the user experience, which was the main aim of the research question (Section 1.1). We have sought to describe the environment, as seen from a user perspective, in a single number so that the geographic extent of problem areas can be represented on a map.

Seen together in graphs such as Figure 27, the available technical measurements are best represented as a joint probability distribution, with the marginal frame rate distribution indicating the utilisation and the marginal retry ratio distribution indicating the level of network degradation.

On the assumption that network degradation is the main issue of interest to users, this section relates the marginal retry ratio distribution to a Mean Opinion Score, which is a suitable parameter for colouring a map display.

8.1 MEAN OPINION SCORE

In order to quantify the user experience, it is possible to use a Mean Opinion Score (MOS) scale. A MOS is a subjective user rating scale that is used for assessing the degree of user satisfaction with a particular service. A wide variety of MOS scales exist for use in various applications, but they are most widely used in audio and video engineering (e.g. Awoniyi and Tobagi, 2004).

By way of an example, consider the voice quality performance of codecs that might be used in a VoIP application. Figure 32 shows the performance that is to be expected from three such codecs as graphs of MOS versus packet loss rate (VoIPTroubleshooter, SA) In this case the recommended range of operation is with a MOS of over 3.5. From the graphs one can see the maximum packet loss rate that can be tolerated to achieve such a quality of service.

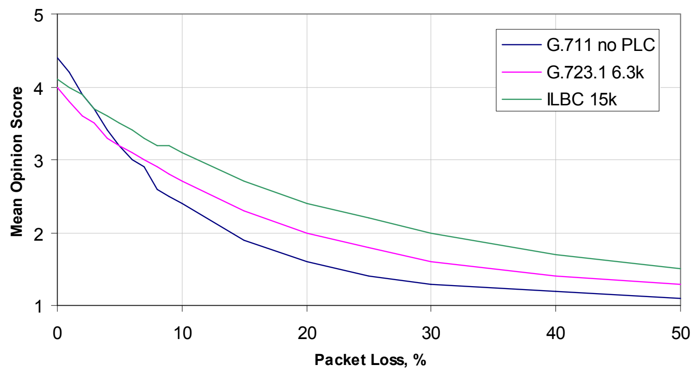

Figure 32 Example of a MOS scale (codec performance)

Typically a MOS scale ranges from 1 to 5 as shown in the table below. There are various versions of this for different types of service[5], but the generic concept of a MOS is useful in that it enables a link to be made between a technical measurement and user experience.

MOS	Quality	Impairment
5	Excellent	Imperceptible
4	Good	Perceptible but not annoying
3	Fair	Slightly annoying
2	Poor	Annoying
1	Bad	Very annoying

In section 6.2 the experimental method was described. This involved measuring the ratio of two stopwatches for timing the amount of time for which video was available over a WiFi link. Figure 33 summarises the relationship between the MOS and the stopwatch ratio (S2/S1). There is generally good agreement between these parameters with a correlation coefficient of 0.95. Hence, it was felt reasonable to use the experimental method of section 6.2 to infer the MOS that a user might expect in a similar environment.

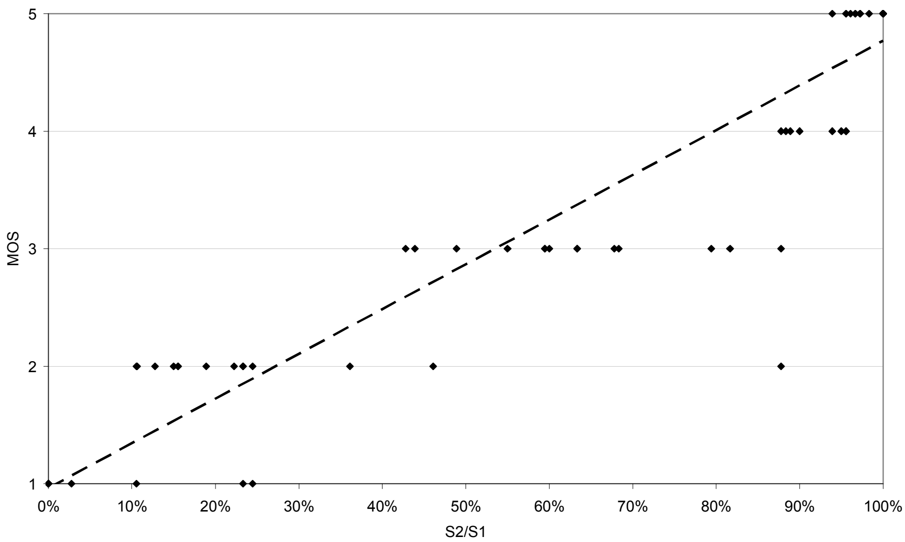

Figure 33 Experimental relationship between MOS and stopwatch ratio

Having arrived at a measure of user experience that seemed like a good basis for answering the research question (section 1.1) it was not then possible to find any correlations in the experimental data that were acceptable. Correlations were unsuccessfully investigated between the MOS and different statistics describing the frame rate, retry rate, retry ratio and bytes/second.

[5] e.g. ITU-T Recommendation P.800 (1996) for voice, ITU-R Recommendation BT.500 (2000) for video

After trying many combinations of parameters it was decided that the MOS could not be inferred directly by any of the passive measurements. This was because there were too many variables affecting each measurement and these variables could not be measured using a passive monitor.

It was concluded that a lower bound on the MOS could, however, be estimated and this approach leads to a useful description of the environment.

8.2 MOSLB

Figure 34 shows all the experimental data (with and without interference) in terms of the MOS versus the mean retry ratio, γ. A MOS of 5 was observed for mean retry ratios as high as 35%. At the other end of the scale, a MOS of 1 was observed for mean retry ratios as low as 3.6%. It is very clear that there is no direct correlation between these parameters.

The only firm observations that can be made from this graph are:

- A MOS of 5 is significantly more likely to occur at low values of mean retry ratio (as shown by the higher density of MOS 5 data points below a mean retry ratio of 10%);
- At very low values of mean retry ratio (below about 10%) it is unlikely that the MOS will be low.

Hence it is concluded that passive measurements can only be used to indicate the lower bound on the MOS. The proposed MOSLB is, therefore, an indication of the worst user experience that is likely to be encountered.

We define MOSLB as the lowest value of MOS that is to be expected at a measured mean retry ratio. By plotting MOS against mean retry ratio in Figure 34 we can define a function for MOSLB to be used in colouring the maps in Appendix C.

The MOSLB proposed is a function of the mean retry ratio and is based on the estimate of the mean retry ratio, $\langle\gamma\rangle$.

$$
\begin{array}{ll}
0\% \le \langle\gamma\rangle < 2.5\% & MOSLB = 5 \\
2.5\% \le \langle\gamma\rangle < 3.5\% & MOSLB = 4 \\
3.5\% \le \langle\gamma\rangle < 5\% & MOSLB = 3 \\
5\% \le \langle\gamma\rangle < 9\% & MOSLB = 2 \\
9\% \le \langle\gamma\rangle & MOSLB = 1
\end{array}
\qquad (2)
$$

This function is shown as a purple line on both Figure 34 and Figure 35. The position for the MOSLB has been estimated by visually placing the line so that it looks reasonable on both graphs, taking into account experimental error and the subjective nature of the MOS.

The MOSLB metric has been used for colouring the maps given in Appendix C. The results agree well with the anecdotal evidence and the results of our tests and observations in different locations.

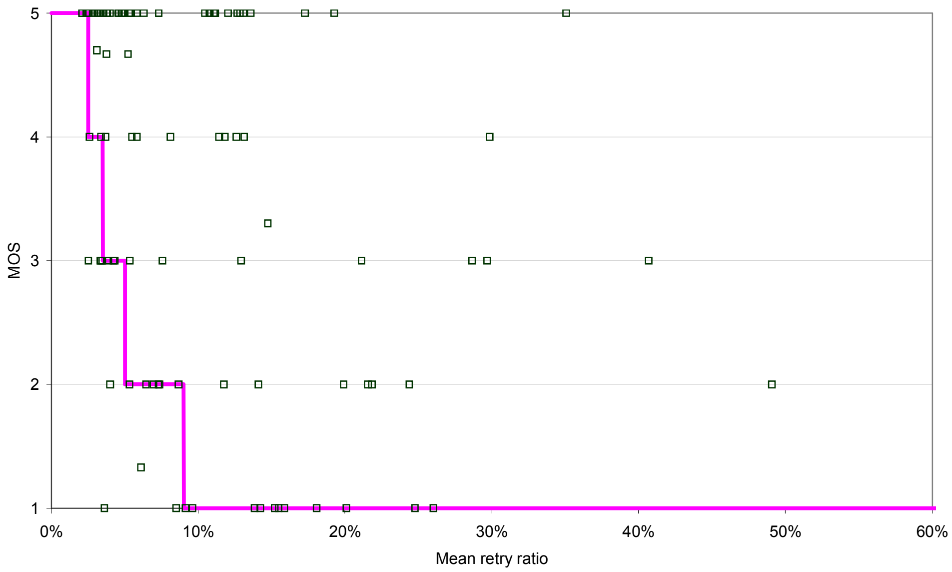

Figure 34 Definition of MOSLB as a function of mean retry ratio

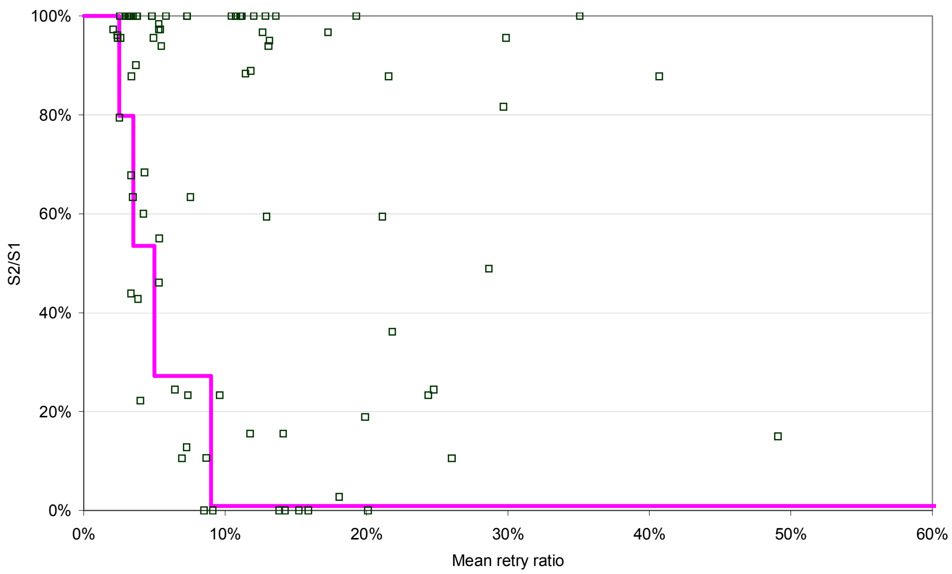

Figure 35 Definition of MOSLB in terms of stop watch ratio versus mean retry ratio

8.3 NETWORK DENSITY

In Section 7.1 it was pointed out that network density is a relatively straightforward measurement to make. People can relate to this measurement so it is often used to describe how busy a site is. Unfortunately the correlation between utilisation and network density is not particularly good (Section 7.3), as not all APs are fully loaded all the time.

In other words, whilst there may be a lot of APs on the list of available networks at a site, it doesn't necessarily follow that the site will be very busy. The network density is such a popular measurement, however, that our results have to be related to it.

Figure 36 shows the MOSLB plotted against the network density (BSSIDs/ch/km^2) for all of the grid squares in our survey. There is not a good correlation between these parameters, implying that the user experience is not directly related to the density of networks in an area, although there is evidence of network degradation at very high network densities.

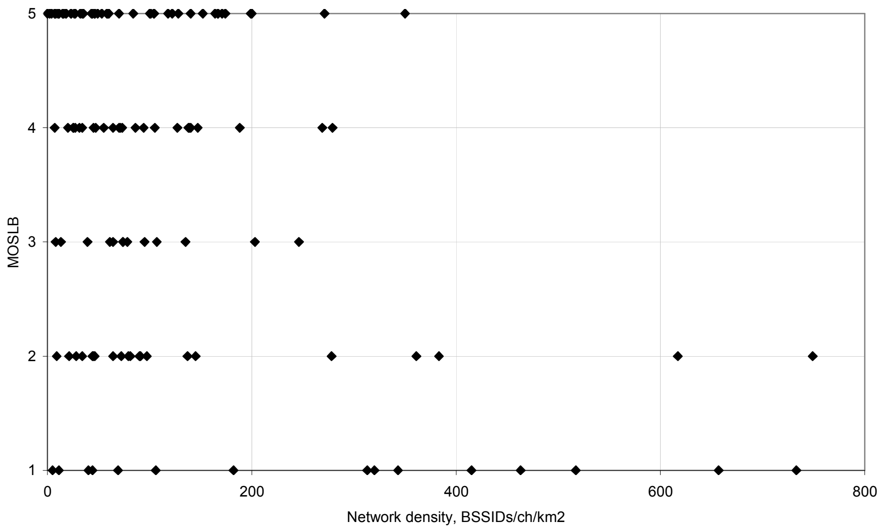

Figure 36 MOSLB versus network density

That this is the case can be seen by some good examples in Appendix C. Considering the cases where the MOSLB is 1, compare the following grid squares, ordered by decreasing network density:

Grid Square	Location	BSSIDs/ch/km^2
TQ 33 81	London, Liverpool Street	657
NJ 94 06	Aberdeen, city centre	106
SO 94 21	Cheltenham	69
TL 02 22	Luton, White Lion retail park	44
TL 43 58	Cambridge	40
TL 22 70	Huntingdon, bar/restaurant	5

A poor MOSLB can be found at any location, but we do not see a MOSLB of 5 at very high network densities (above 400 BSSIDs/ch/km^2).

8.4 CHANNEL BALANCE

One of the anecdotal reports (section 2.2.1) was that network problems are caused by people using one channel more than the others, thereby placing an excessive load on part of the spectrum.

To investigate this conjecture a channel balance parameter, β, has been defined as:

$$\beta = \left(1 - \sqrt{\frac{N}{N-1}\sum_{n=1}^{N}\left(\frac{m}{M} - \frac{1}{N}\right)^2} \right).100\% \qquad (3)$$

where N is the number of channels (in this case 3, as channels 1, 6 and 11 were monitored), m is the number of BSSIDs on channel n and M is the total number of BSSIDs.

With this definition, then, if β = 0, only one channel is in use. If β = 100% then all the channels are being used equally.

Figure 37 is a plot of the MOSLB for each grid square surveyed against the channel balance.

If the channel balance is above 85% then the MOSLB increases, suggesting that a good user experience is likely if all the channels are equally used. However, a low value of channel balance does not automatically imply a poor MOSLB. In general the level of network degradation is not well correlated with the channel balance.

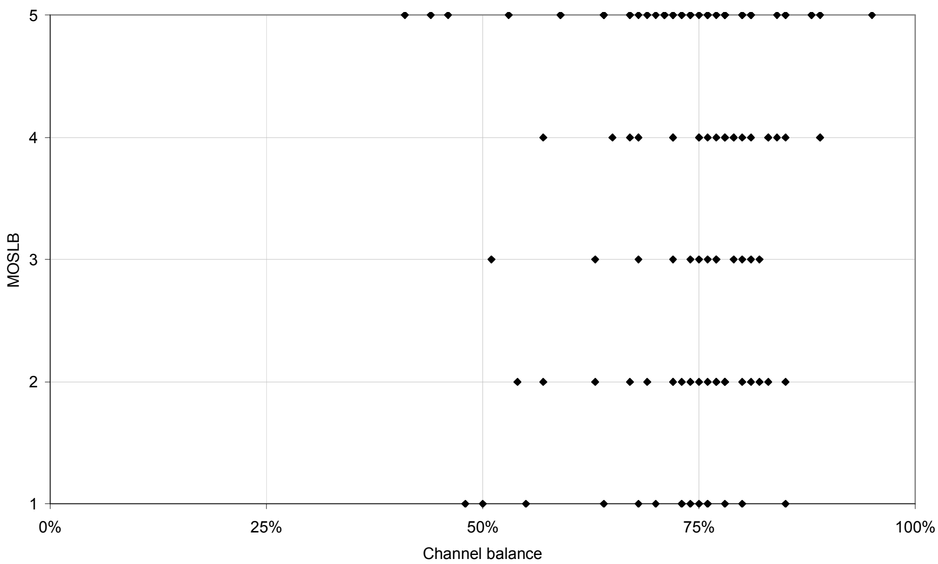

Figure 37 MOSLB versus channel balance

Figure 38 shows the channel balance of each grid square plotted against the network density.

The lower values of channel balance tend to occur at the lowest network densities, which is to be expected. At network densities above 200 BSSIDs/ch/km^2 the channel balance is above 60%, suggesting that in busy, urban locations the channels are reasonably well balanced.

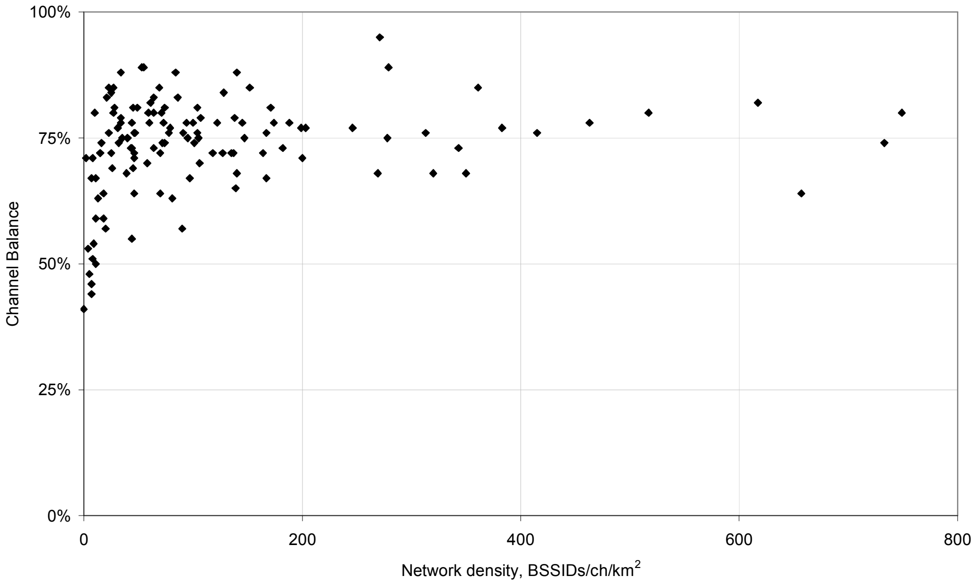

Figure 38 Channel balance versus network density

8.5 SUMMARY

This section has expanded the analysis of the survey and experimental data to arrive at an indication of worst-case user experience that would be expected in an area.

The metric we have derived for quantifying the user experience is the MOSLB which is derived from the mean retry ratio. The MOSLB has been used to colour code the maps in Appendix C and provides evidence for the problems reported in the centre of London.

Channel balance does not appear to be a significant problem. In the busier areas there is an even spread of usage between channels 1, 6 and 11.

9 FUTURE TRENDS

Whilst the main interest of this project has been in measuring the usage of the WiFi networks, various factors affecting the future usage have become apparent. These are listed here as they are likely to affect any policy decisions concerning the LE bands.

There are a number of trends that are likely to impact on the future use of the 2.4 GHz ISM band. The bands is already clearly stressed in central London due to interfering signals, so any new services may struggle to provide the user experience that the vendors would desire.

1 **Media players**. These devices are likely to increase the load placed on WiFi networks in residential areas. Companies such as D-Link (www.dlink.com) are producing media players for the domestic market that integrate a number of media streaming applications into one box. These typically have high-definition HDMI video outputs and IEEE 802.11g connectivity. They allow users to stream audio, images and video from PC to television as well as storing content. With a media player a home user can search video web sites (e.g. www.youtube.com) using their television instead of a PC.

2 **Voice over IP (VoIP)**. This is already a well-developed technology and VoIP is now included in multimode handsets. There is a concern that VoIP services may become unworkable in areas where the WiFi network degradation is too severe.

3 **Bluetooth over IEEE 802.11**. A version of Bluetooth is being developed that runs on the IEEE 802.11 MAC/PHY layers in order to increase the data rate. This move is intended to extend the services available via Bluetooth before Ultra Wide Band (UWB) solutions become available. If the move to IEEE 802.11 is approved then chips are likely to start appearing in 2009/10 (Wilson, 2008). Possible applications for this technology would be bulk transfers of music photos and videos between PCs, cameras, phones, etc.

4 **Urban communications**. The use of urban communications is now widespread, with different spectrum bands and technologies being used for different applications. As an example, NOW Wireless (www.nowwireless.com) has been deploying urban networks for over five years and use a mixture of communications services including GSM, 3G and a 2.4 GHz/5GHz system called Mesh4G (NOW Wireless, 2007). Using LE bands is clearly cheaper for councils, so it is reasonable to expect increased reliance on these bands for urban street furniture communications. This finding supports that of Ofcom (2008b) which predicted an increase in reliance on the 2.4 GHz and 5 GHz bands in coming years.

5 **IEEE 802.11a**. The use of IEEE 802.11a in the 5 GHz band is not yet as widespread as WiFi in the 2.4 GHz band. As more devices appear in the 5 GHz band, some users will move across to that band, especially if they get better quality of service by doing so.

6 **RF Barrier**. With WiFi eavesdropping being a concern for many users, technologies are appearing to thwart potential attacks. One such technology is the RF Barrier technology (Saran, 2008) which is an active system designed to jam WiFi receivers outside a building from detecting signals emanating from within the building.

7 **Sensor networks**. A wide variety of sensor networks is appearing that use the LE bands. Typically such networks only require small amounts of bandwidth and use protocols such as ZigBee and Bluetooth rather than WiFi. The concern for these systems would be the impact of increased network degradation.

8 **Channel bonding**. Using two channels to increase the available data rate is attractive for developers trying to improve their products and would facilitate higher bandwidth services. There are concerns that this technology, which is part of the IEEE 802.11n standard, will cause interference with existing IEEE 802.11b/g networks, so channel bonding is not currently being implemented at 2.4 GHz.

9 **Roaming**. The original IEEE 802.11 standards did not allow for roaming between APs. The 802.11f and 802.11r standards are now ratified and will improve the ability for clients to roam between APs and assist with load balancing. IEEE 802.11r will be particularly important to VoIP phone users, who will see improved call continuity in enterprise applications. This should facilitate the use of corporate WiFi networks for VoIP telephony.

10 **Wireless network management**. IEEE 802.11v builds on the information provided by IEEE 802.11k to enable clients to be managed remotely. Once the infrastructure and services are available to support network management the client devices should see improved battery life and better synchronisation for multimedia streaming.

11 **Digital AV senders**. The current generation of analogue AV senders is being replaced in the baby monitoring market and wireless security camera market by digital alternatives. There is no sign of this change in the domestic TV relay market, probably because the price increase would be too high for the benefits. It is reasonable to assume, however, that WiFi-friendliness and functionality provided by digital modulations, will eventually lead to all these markets moving to digital devices.

The 2.4 GHz ISM band is clearly attractive to many developers, as it promises to provide low-cost, high-bandwidth networking over a useful range. The success of the WiFi family of products and services clearly demonstrates this.

The success of future technologies in this band is dependent on interoperability between the protocols. Our results suggest that interference between existing systems is already a significant problem and that many systems are already deployed that have not been designed to be interoperable.

Some technology developments will help to overcome the co-existence problems. One of these will be the move away from analogue AV senders to digital alternatives. Another will be the automatic configuration of communications protocol stacks, with less reliance on users to configure their devices.

10 CONCLUSIONS

A number of conclusions have been reached during this project. They relate to the problems that users are experiencing, the metric for measuring the user experience and the practicality of the monitoring method. Also, thoughts on the future pressures on the 2.4 GHz ISM band are presented. Finally, suggestions are given on future work in this area.

10.1 WIFI PROBLEMS

1 Between 10% and 15% of users experience problems when using WiFi access to the Internet. This is an average and may well be substantially higher in the inner city areas.

2 Users tend to attribute their problems to congestion, even if their problems are caused by factors other than congestion. The majority of their problems do not stem from the wireless networks at all, but are due to a plethora of other causes, including device configuration and wired network problems.

3 Where problems do reside in the wireless networks, they are more often due to interference than congestion.

4 The level of WiFi network degradation is not evenly spread across the country and is much higher inside the major cities than outside. In central London, there are some locations where congestion is likely to be occurring regularly. The railway station areas seem to be particularly prone to this, as they have very high network densities. We expect a similar situation to occur in each of the major cities, but more survey data is needed to confirm this.

5 The density of WiFi networks in an area is not, *per se*, a good indicator of potential WiFi degradation. However, a high density of WiFi networks implies a high population density and this implies an increased probability of encountering interference and other network degradation effects.

6 Outside the cities network degradation is not associated with high population density. It tends to be very localised (<1km) and is therefore usually attributable to isolated interference events.

7 The information presented to users by their devices leads to the idea that problems are due to the wireless networks being too busy and the word 'congestion' is consequently used to describe any networking problem. Our results do not support this concept, rather they suggest that the IEEE 802.11 protocols can support much higher loads than are seen in practice. The bulk of the problems seen by users are not, therefore, due to excessive utilisation, rather they can be attributed to a variety of effects including: interference from analogue devices (e.g. AV senders, microwave ovens), device configuration errors, wired network problems, etc.

8 Residential users are often blamed for causing problems by excessive use of their WiFi connections, but no evidence has been found to substantiate this. It is felt that the bulk of such problems can be attributed to the use of AV senders, which are particularly harmful to WiFi signals. Such interference incidents will be very localised.

10.2 MEASUREMENT METRIC

1 We have investigated a metric based on the Mean Opinion Score (MOS) concept and have found that there is no simple relationship between any of the parameters that can be measured passively and the MOS. We have, however, proposed that the lower bound on the MOS (referred to here as the MOSLB) can be estimated and can be used to identify areas that might be suffering network degradation at the physical or wireless link layer.

2 Of all the parameters we have explored, the most useful parameter that can be measured to infer the MOSLB is the ratio of the retry rate to the total frame rate. This parameter can be measured reliably and easily using handheld devices.

3 The MOSLB metric allows the wireless part of the WiFi access to be observed. It does not depend on the performance of the wired network and is therefore a useful indicator of wireless network degradation.

10.3 PRACTICALITY OF THE MONITORING METHOD

1 Passive monitoring of the 2.4 GHz ISM band using small, handheld receivers is practical and, because of the short ranges of communications links at this frequency and difficulty of vehicular access in urban areas, is best carried out by walking or cycling around the areas of interest. A mixture of *ad hoc* and targeted monitoring would provide good coverage of the UK.

2 A resolution of 1 km is appropriate and a minimum of 15 minutes recording should be carried out in every 1 km grid square in order to obtain useful statistics on network behaviour. Longer periods of recording yield better statistics and are to be encouraged. A practical target would be five grid squares per day.

3 Using a built-in GPS receiver to track the device is an effective approach even if the accuracy is rather limited. If a 1 km resolution is used, then the accuracy of such receivers is adequate and the only significant error source is loss of the GPS signals in tunnels and indoors. The London Underground should be avoided and GPS lock should be obtained before entering buildings.

10.4 FUTURE DEMAND

1 We anticipate that the pressures on the 2.4 GHz band will carry on increasing in the near future and that 5 GHz solutions will be deployed to meet some of this demand. However, many of the pressures on the 2.4 GHz band are not due to increases in the offered load, but to the 'wastage' of the available spectrum through interference, greedy usage, etc.

2 High-definition television is not likely to be deployed extensively in the 2.4 GHz band. If the Wireless High Definition Interface (WHDI) standard becomes popular, however, that will increase the load on the 5 GHz Unlicensed National Information Infrastructure (UNII) band, meaning that the 2.4 GHz ISM band will remain the dominant band for Wireless Local Area Networks (WLAN).

3 The lack of licence costs makes the LE bands attractive to councils running urban networks. Such networks will have to be very resilient to deal with the high levels of network degradation that have already been seen in central London.

10.5 FURTHER WORK

The survey work carried out in this project has proved the principle of monitoring in urban areas using small, handheld receivers and has led to the development of a performance metric.

We recommend repeating the survey every two years to look for trends in utilisation and network degradation. It may be necessary to reassess the resolution at which the results are compared once such future data becomes available.

We have only looked at one type of AV sender in this project. It would be worth investigating other types of AV senders, including wireless security cameras and baby monitors, to see whether they all exhibit the same kind of effects when used in proximity to WiFi networks. Comparison with the newer digital systems would also be useful, especially as there are two different types of digital modulation appearing in products now on sale.

The survey has shown that the centre of London exhibits extensive network degradation. As a next step, it is proposed that the survey be extended further out from the centre of London to gauge the geographical extent of the degraded region. Furthermore, it is proposed that the same kind of survey is performed in Birmingham and another major UK city to ascertain whether the same effects can generally be expected in the large cities.

Another useful extension to the survey would be to consider diurnal variations in utilisation and network degradation. Previous studies using stationary receivers have shown considerable changes in the levels of usage throughout the day and such information can be useful in predicting the demand for services at different times in different types of location. It is suggested that map displays, similar to those given in Appendix C of this report, would be helpful in visualising such patterns of behaviour.

It would also be useful to consider ways in which the use of the 2.4 GHz band can be improved to make better use of the available spectrum. The idea of a '2.4 GHz friendly' accreditation has been proposed after considering the evidence of this work. How such a scheme could be made to work in practise and which radio systems are most likely to be affected, would be a worthwhile area for investigation.

11 ABBREVIATIONS

ACR	Absolute Category Rating
AIMS	Autonomous Interference Monitoring System
AP	Access Point
AV	Audio Video
BSSID	Basic Service Set Identifier
BNG	British National Grid
CCR	Comparison Category Rating
DCR	Degradation Category Rating
DFS	Dynamic Frequency Selection
DHCP	Dynamic Host Control Protocol
DIY	Do It Yourself
DMOS	Differential Mean Opinion Score
DOS	Denial Of Service
GPS	Global Positioning System
GSM	Global System for Mobile communications
HDMI	High Definition Multimedia Interface
HTTP	HyperText Transfer Protocol
IP	Internet Protocol
ISM	Industrial, Scientific & Medical
ISP	Internet Service Provider
IT	Information Technology
LE	Licence Exempt
LLC	Logical Link Control
MAC	Media Access Control
MOS	Mean Opinion Score
MOSLB	Mean Opinion Score Lower Bound
PHY	Physical layer
QoS	Quality of Service
RF	Radio Frequency
RSGB	Radio Society of Great Britain
RTS/CTS	Request To Send/ Clear To Send
SA	Sine Anno
SSID	Service Set Identifier
TCP	Transmission Control Protocol
TV	Television
UNII	Unlicensed National Information Infrastructure
UWB	Ultra Wide Band
VoD	Video on Demand
VoIP	Voice over Internet Protocol
VRRP	Virtual Router Redundancy Protocol
VSAQ	Video Service Audio Quality
WHDI	Wireless High Definition Interface
WiGLE	Wireless Geographic Logging Engine
WLAN	Wireless Local Area Networks

12 DEFINITION OF TERMS

Congestion	A network link is congested when the offered load on the link reaches a value close to the capacity of the link (Jardosh *et al*, 2005a)
Degradation	A network is suffering degradation when it cannot provide the maximum performance to the users. Degradation effects include congestion, interference and all the other problems listed in section 4.2.
Interference	Noise or other external signals coming from other devices, such as microwave ovens and other wireless network devices that will result in delay to the user either by blocking transmissions from stations on the WLAN or by causing bit errors to occur in data being sent (Vines, 2002, pp. 234-235)
Network density	The network density is quoted in this report in terms of the number of unique BSSIDs per channel per square kilometre.
Offered load	The total number of bits presented for transmission over a wireless network per second.
Throughput	The throughput is the total number of bits transmitted over a wireless channel per second (Jardosh *et al*, 2005a). As not all the offered load may be carried successfully by the network, throughput can be less than offered load.

13 REFERENCES

Aberdeen City Council (2008) *Population report, Aberdeen city and shire.* Available from: http://www.aberdeencity.gov.uk/nmsruntime/saveasdialog.asp?lID=17218&sID=3365 (Accessed 5 January 2009)

adslnation (2009) *ADSL microfilters*, [online] Available from: http://www.adslnation.com/support/filters.php (Accessed 1 January 2009)

Anastasi, G., Borgia, E., Conti, M. and Gregori, E. (2004) 'Wi-Fi in ad hoc mode: a measurement study' *in Proceedings of the IEEE International Conference on Pervasive Computing and Communications*, Orlando, Florida, March 14-17, 2004.

Awoniyi, O. and Tobagi, F.A. (2004) 'Effect of fading on the performance of VoIP in IEEE 802.11a WLANs'. *2004 IEEE Conference on Communications*, vol. 6, 20-24 June 2004, pp. 3712-13717 [online] Available from: http://mmnetworks.stanford.edu/papers/Awoniyi_icc04.pdf (Accessed 12 January 2009)

Barsocchi, P., Oligeri, G. and Potorti, F. (2005) *Validation for 802.11b wireless channel measurements*, [online] Available from: http://wnlab.isti.cnr.it/gabriele/papers/Vfor80211bWCM/Vfor80211bWCM.pdf (Accessed 9 January 2009)

Battiti, R., Conti, M., Gregori, E. and Sabel, M. (2003) 'Price-based congestion-control in Wi-Fi hot spots'. *in Proc. WiOpt'03* [online] Available from: http://rtm.science.unitn.it/~battiti/archive/Battiti26-wiopt.pdf (Accessed 9 January 2009)

Brik, V., Rayanchu, S., Saha, S., Sen, S., Shrivastava, V. and Banerjee, S. (2008) *A measurement study of a commercial-grade urban WiFi mesh*, [online], IMC '08, October 20-22, 2008 Available from: http://pages.cs.wisc.edu/~suman/pubs/mcb.pdf (Accessed 15 September 2008)

Bruce, W.R. (2002) *Wireless LANs end to end*. New York, Hungry Minds Inc.

Chatzimisios, P., Boucouvalas, A.C. and Vitsas, V. (2005). 'IEEE 802.11 wireless LANs: Performance analysis and protocol refinement' *EURASIP Journal on Applied Signal Processing* 2005:1, 67-78

Cisco (SA) *20 myths of Wi-F interference: Dispel myths to gain high-performing and reliable wireless.* CISCO White Paper. [online] Available from: http://www.cisco.com/en/US/prod/collateral/wireless/ps9391/ps9393/ps9394/prod_white_paper0900aecd807395a9.pdf (Accessed 1 January 2009)

Cochrane, P. (2007) *Wi-fi* London, [online] Available from: http://networks.silicon.com/mobile/0,39024665,39165548,00.htm (Accessed 5 January 2009)

Cox, J. (2007) *Strange Xbox signal suspected of jamming wireless LANs*, [online] Available from: http://www.networkworld.com/news/2007/121307-microsoft-xbox-jams-wireless-lans.html (Accessed 1 January 2009)

Cunningham, S. and Grout, V. (2007) 'War and peace: a practical study of Wi-Fi related issues'. *Proceedings of the International Conference E-Activity and Leading Technologies (E-ALT '07),* Porto (Oporto), Portugal, 3-6 December 2007. [online] Available from: http://www.wrexham.ac.uk/cunninghams/research/warandpeace.pdf (Accessed 16 September 2008)

Day, S.W. and Merricks, N.P. (2003) *AY4434 – 2.4GHz Monitoring Exercise*. Mass Consultants Ltd. MC/SC0390/REP010/1

De Maesschalck, T. (2006) *Wireless router trouble shooting guide*, [online] Available from: http://www.dvhardware.net/article15009.html (Accessed 1 January 2009)

Evans, D. (2008) *iPlayer piling the pressure on worried ISPs*. [online], Computing, 1 May 2008. Available from: http://www.computing.co.uk/computing/analysis/2215639/iplayer-piling-pressure-worried-3967566 (Accessed 10 September 2008)

Franceschinis, M., Mellia, M., Meo, M., and Munafo, M (2005) 'Measuring TCP over WiFi : A real case'. *1st Workshop on Wireless Network Measurements,* Riva del Garda, Italy, 3 April 2005. [online] Available from: http://www.tlc-networks.polito.it/mellia/papers/winmee.pdf (Accessed 9 January 2009)

Frary, M. (2008) *Infected by the free public WiFi virus*, [online], Times Online. 14 March 2008. Available from: http://www.timesonline.co.uk/tol/travel/business/article3554838.ece (Accessed 31 October 2008)

Freenet (2008) *Hotel Warrior Kit*. [online] Available from: http://store.freenet-antennas.com/product_info.php?cPath=22_54&products_id=233 (Accessed 15 September 2008)

Gardner, W.D. (2003) *Urban Wi-Fi- Gridlock Predicted to Arrive in 2004*, [online] Available from: http://www.techweb.com/wire/story/TWB20031024S0011 (Accessed 15 September 2008)

Geier, J. (2002) *Improving WLAN performance with RTS/CTS*. [online] Available from: http://www.wi-fiplanet.com/tutorials/article.php/1445641 (Accessed 1 January 2009)

Golmie, N. (SA) *Interference in the 2.4 GHz ISM band: Challenges and solutions* [online] Available from http://www.antd.nist.gov/pubs/golmie.pdf (Accessed 8 January 2009)

Goodwin, M. (2007)*10 tips to increase the range of your wireless* network, [online] Available from: http://www.dailywireless.com/features/10-tips-increase-wireless-range-051007/ (Accessed 9 September 2008)

Gubbins, E. (2003) W*hen Will Wi-Fi Make Some Dough,* [online] Available from: http://telephonyonline.com/wireless/mag/wireless_wifi_dough/ (Accessed 15 September 2008)

Hansell, P., Kirtay, S., Inglis, I., Pahl, J. and Munday, S. (2004) *Evaluating spectrum percentage occupancy in licence-exempt allocations*. Aegis Systems Ltd. and Transfinite Systems Ltd. 1606/LEM/R/3

Heusse, M., Rousseau, F., Berger-Sabbatel, G. and Duda, A. (2003) 'Performance anomaly of 802.11b'. *INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies,* Vol. 2, pp. 836-843

Higgins, T. (2007) *Hole discovered in Wi-Fi 802.11n draft 2.0 certification* test, [online] Available from: http://www.smallnetbuilder.com/content/view/30196/97/ (Accessed 1 January 2009)

Horvitz, R. (2008) *Interference issues: Is the 2.4-GHz band reaching supersaturation?*, [online] Available from: http://w2i.com/resource_center/the_w2i_report__weekly_newsletter/news/p/newsletterId_175/id_213 (Accessed 10 September 2008)

IETF (1989) *Request for Comments: 1122. Requirement for Internet Hosts -- Communications Layers*. [online] Available from: http://tools.ietf.org/html/rfc1122 (Accessed 24 March 2009)

IEEE (2007) *IEEE Standard for Information Technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY)* Specifications, [online] Available from: http://standards.ieee.org/getieee802/download/802.11-2007.pdf (Accessed 7 January 2008)

informitv (2008) *BBC iPlayer usage prompts concern from service providers*, [online] Available from: http://informitv.com/articles/2008/02/20/bbciplayerusage/ (Accessed 7 January 2009)

Intel (2008) *What is channel bonding?*, [online] Available from: http://www.intel.com/support/wireless/sb/CS-025343.htm (Accessed 2 February 2009)

Intellect (2008) *Capitalising on Convergence 2: The little red light.* [online] Available from: http://www.intellectuk.org/component/option,com_docman/task,doc_download/gid,2388/Itemid,102/ (Accessed 10 September 2008)

ITU-R Recommendation BT.500 (2000) *Methodology for the subjective assessment of the quality of television pictures.* International Telecommunications Union, Geneva, Switzerland

ITU-T Recommendation P.800 (1996) *Methods for subjective determination of transmission quality*. International Telecommunications Union, Geneva, Switzerland

Jardosh, A.P., Ramachandran, K.N., Almeroth, K.C. and Belding-Royer, E.M. (2005) 'Understanding congestion in IEEE 802.11b wireless networks', in *Proceedings of the Internet Measurement Conference 2005,* Berkeley, California, USA, pp. 279-292

Jardosh, A.P., Ramachandran, K.N., Almeroth, K.C. and Belding-Royer, E.M. (2005) 'Understanding link-layer behavior in highly congested IEEE 802.11b wireless networks' *SIGCOMM '05 Workshops*, 22-26 August 2005, Philadelphia, USA.

Jun, J., Peddabachagari, P. and Sichitiu, M. (2003) 'Theoretical maximum throughput of IEEE 802.11 and its Applications'. *Second International Conference on Network Computing and its Applications*, 16-18 April, 2003, pp. 249-256.

Kaleem, Z. (2007) "*Free Public WiFi" SSID.* [online] Available from: http://www.wlanbook.com/free-public-wifi-ssid/ (Accessed 5 January 2009)

Kemerlis, V.P., Stefanis, E.C., Xylomenos, G. and Polyzos, G.C. (2006) 'Throughput unfairness in TCP over WiFi'. *in Proc. WONS 2006*. [online] Available from: http://hal.archives-ouvertes.fr/docs/00/05/46/08/PDF/04-Kemerlis.pdf (Accessed 9 January 2009)

Kershaw, M. (2004) *802.11 networking in hostile environments.* [online] Available from: http://www.kismetwireless.net/presentations/mit-hostilewireless.pdf (Accessed 20 January 2009)

Kewney, G. (2003) *Breaking WiFi Gridlock*, [online] Available from: http://www.eweek.com/c/a/Mobile-and-Wireless/Breaking-WiFi-Gridlock (Accessed 15 September 2008)

Leeson, H., Hansell, P., Burns, J. and Spasojeviç,Z. (2000) *Demand for use of the 2.4 GHz ISM band. Final report.* Aegis Systems Ltd. 1215/AE/ISM2/R/2

Li, N., Yan, B. and Chen, G. (2008) *A measurement study on wireless camera networks*. [online] Available from: http://www.cs.uml.edu/~glchen/papers/camera-icdsc08.pdf (Accessed 25 March 2009)

Linksys (2009) *Wireless router configuration – advanced wireless settings.* [online] Available from: http://www1.linksys.com/support/troubleshoot/wireless/ wireless_router_config/advanced_settings.html (Accessed 1 January 2009)

Macdonald, N. (2004) *Wi-Fi in the real world*, [online] Available from: http://www.theregister.co.uk/2004/02/10/wifi_in_the_real_world/ (Accessed 1 January 2009)

Mac-net (SA) *Free public WiFi virus.* [online] Available from: http://www.mac-net.com/1533485.page (Accessed 7 January 2009)

Mahanti, A., Williamson, C. and Arlitt, M. (2007) 'Remote analysis of a distributed WLAN using passive wireless-side measurement'. *Performance Evaluation*, 64(2007) pp. 909-932.

McCullough, J. (2004). *185 Wireless Secrets*, Indianapolis, USA, Wiley Publishing Inc.

Mitchell, B. (SA) *How to build a wireless home network – tutorial*. [online] Available from: http://compnetworking.about.com/cs/wirelessproducts/a/howtobuildwlan.htm (Accessed 1 January 2009)

Molta, D. (2003) *Spreading Wi-Fi Rumours*, [online] Available from: http://www.networkcomputing.com/showitem.jhtml?docid=1423buzz4 (Accessed 16 September 2008)

NOW Wireless (2007) *Wireless Traffic Lights* [online] Available from: http://www.nowwireless.com/metro/pdf/wp/wireless_trafficlights_wp_v071026.pdf (Accessed 9 January 2009)

O'Connor, P.D.T. (1981) *Practical Reliability Engineering*. London, Heyden & Son Ltd.

Ofcom (2007). *Licence-Exemption Framework Review: A statement on the framework for managing spectrum used by licence-exempt devices,* [online] Available from: http://www.ofcom.org.uk/consult/condocs/lefr/lefr_statement/lefr_statement.pdf (Accessed 10 September 2008)

Ofcom (2008) *UK broadband speeds 2008.* [online] Available from: http://www.ofcom.org.uk/research/telecoms/reports/bbspeed_jan09/bbspeed_jan09.pdf (Accessed 12 January 2009)

Ofcom (2008) *Tomorrow's Wireless World.* [online] Available from: http://www.ofcom.org.uk/research/technology/overview/randd0708/#_ftn1 (Accessed 12 January 2009)

Office for National Statistics (2001) *KS01 Usual resident population: Census 2001, Key Statistics for urban areas.* Available from: http://www.statistics.gov.uk/StatBase/ssdataset.asp?vlnk=8271&Pos=2&ColRank=1&Rank=224 (Accessed 3 June 2008)

Pätzold, M. (2002) *Mobile fading channels.* New York, USA, John Wiley & Sons Inc.

Petrin, A.J. (2005) *Maximizing the utility of radio spectrum: Broadband spectrum measurements and occupancy model for use by cognitive radio.* PhD Thesis, Georgia Institute of Technology. Available from: http://etd.gatech.edu/theses/available/etd-07152005-135311/unrestricted/petrin_allen_j_200508_phd.pdf (Accessed 10 September 2008)

Phifer, L. (2004) *Settings for RTS threshold and fragmentation threshold,* [online] Available from: http://searchmobilecomputing.techtarget.com/expert/KnowledgebaseAnswer/0,289625,sid40_gci1009360,00.html (Accessed 1 January 2009)

Phifer, L. (2007) *Wireless network troubleshooting: Connectivity*. [online] Available from: http://searchnetworking.techtarget.com/news/article/0,289142,sid7_gci945257,00.html# (Accessed 1 January 2009)

Purdy, K. (2007) *Boost your WiFi signal for less than a dollar*. [online] Available from: http://www.lifehacker.com.au/tips/2007/11/21/boost_your_wifi_signal_for_les.html (Accessed 12 January 2009)

Putman, B.W. (2005) *802.11 WLAN hands-on analysis,* AuthorHouse, [online] Available from: http://books.google.co.uk/books?id=8wvfAw6_Tq0C&pg=PA28&dq=putman+wifi (Accessed 31 October 2008)

Raghavendra, R., Belding, E.M., Papagiannaki, K. and Almeroth. K.C. (2007) *Understanding handoffs in large IEEE 802.11 wireless networks*, [online], IMC '07 October 24-26, 2007, San Diego, California, USA. Available from: http://www.imconf.net/imc-2007/papers/imc192.pdf (Accessed 15 September 2008)

Richardson, K. (2006) 'Polite Protocols'. *Ofcom R&D Symposium*, 23 Nov 2006. Available from: http://www.ofcom.org.uk/research/technology/events/rd2006/polprot.pdf (Accessed 10 September 2008)

Rodrig, M., Reis, C., Mahajan, R., Wetherall, D. and Zahorjan, J. (2005) *Measurement-based characterization of 802.11 in a hotspot setting*, [online], SIGCOMM '05 Workshops, August 22-26 2005, Available from: http://conferences.sigcomm.org/sigcomm/2005/paper-RodRei.pdf (Accessed 7 January 2009)

Srivastwa, R. (2008) *"Free Internet Access" "Free Public WiFi" The viral SSID.* [online] Available from: http://www.clubhack.com/blog/2008/11/08/free-internet-access_viral_ssid/ (Accessed 5 January 2009)

Saran, C. (2008) *Smart RF blocker beats war drivers*. Computer Weekly, 28 October - 3 November 2008, p.16

Sutherland, E. (2003) *Is the WiFi End Nigh?*, [online] Available from: http://www.wi-fiplanet.com/columns/article.php/3104131 (Accessed 7 January 2009)

Tannenbaum, A.S. (1996) *Computer networks*. 3rd ed. New Jersey, Prentice-Hall Inc.

Ven, M.V. (2009) *Speed Tests. 802.11b/g mixed vs. 802.11 g only (with and without G Nitro)*, [online] Available from: http://www.giantmike.com/tests/WirelessRouter.html (Accessed 7 Jan 2009)

VoIPTroubleshooter (SA) *Measuring Voice Quality,* [online] Available from http://www.voiptroubleshooter.com/basics/mosr.html (Accessed 7 January 2009)

Vines, R.D. (2002) *Wireless security essentials.* Indianapolis, Indiana, Wiley Publishing Inc.

Wagstaff, A.J. and Merricks, N.P. (2007). *Autonomous Interference Monitoring System. Phase 2. Final Report.* Mass Consultants Ltd. MC/SC0585/REP016/1

Wellens, M., Wu, J. and Mähönen, P. (2007) *Evaluation of spectrum occupancy in indoor and outdoor scenario in the context of cognitive radio*. [online] Available from: http://www.mobnets.rwth-aachen.de/fileadmin/templates/images/PublicationPdfs/2007/SpectOcc.pdf (Accessed 10 September 2008)

Wilson, R. (2008) *WiFi set to boost Bluetooth data rates*. Electronics Weekly, 29 October - 4 November 2008, p. 8.

X10 (SA) *How do our customers use their video senders?* [online] Available from: http://www.x10europe.com/general/newsletter_he1.htm (Accessed 25 March 2009)

Yeo, J., Youssef, M., Henderson, T. and Agrawala, A. (2005) 'An accurate technique for measuring the wireless side of wireless networks'. *International Conference on Mobile Systems, Applications and Services*. Seattle, Washington, USA. 2005. pp. 13-18

Yoo, C.S. (2006) 'Network neutrality and the economics of congestion'. *Georgetown Law Journal, Vol. 94, June 2006; Vanderbilt Law and Economics Research Paper No. 05-28; Vanderbilt Public Law Research Paper No. 05-33* [online] Available from: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=825669 (Accessed 12 January 2009)

14 BIBLIOGRAPHY

Clark, M.P (2000) *Wireless access networks.* Chichester, West Sussex, John Wiley & Sons Ltd.

Kaj, I. (2002) *Stochastic modeling in broadband communication systems.* Society for Industrial and Applied Mathematics

Noel, T. (2002) *Wireless mobile phone access to the internet.* London, Hermes Penton Ltd.

Sadiku, M.N.O. (2002) *Optical and wireless communications*. Boca Raton, Florida, CRC Press LLC.

